

UCC 5G UPF Configuration and Administration Guide, Release 2026.03

First Published: 2026-07-24

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)

© 2026 Cisco Systems, Inc. All rights reserved.



CONTENTS

Full Cisco Trademarks with Software License ?

PREFACE

About this Guide xxxv

Conventions Used xxxv

CHAPTER 1

5G Architecture 1

Feature Summary and Revision History 1

Summary Data 1

Revision History 1

Architectural Overview 2

Control Plane Network Functions 2

User Plane Network Function 3

Subscriber Microservices Infrastructure Architecture 3

Control Plane Network Function Architecture 4

CHAPTER 2

5G-UPF Overview 7

Feature Summary and Revision History 7

Summary Data 7

Revision History 7

Product Description 8

Use Cases and Features 8

Configuration and Deployment Requirement for UPF 8

Anchor Point for Intra-RAT and Inter-RAT Mobility 9

External PDU Session Point of Interconnect to Data Network 9

Packet Inspection 10

Per QCI packet drop statistics 10

User Plane Part of Policy Rule Enforcement 10
Lawful Intercept 10
Traffic Usage Reporting (Charging) 11
QoS Handling for User Plane 11
Downlink Packet Buffering and Data Notification Triggering 11
Forwarding End Markers to the Source NG-RAN Node 11
MVNO Support 12
Smart Licensing 12
Deployment Architecture and Interfaces 12
UPF Architecture 12
UPF Deployment Architecture 13
VM Deployment 14
Supported Interfaces 17
Standards Compliance 17
Limitations 18

CHAPTER 3 **1:1 Redundancy 19**
Feature Summary and Revision History 19
Summary Data 19
Revision History 19
Feature Description 20
How it Works 20
Configuring 1:1 UPF Redundancy 27
Configuring BFD Monitoring Between Active UPF and Standby UPF 27
Configuring BGP Status Monitoring Between Each UPF and Next-Hop Router 28
Configuring Sx/N4 Monitoring on the Active UPF and Standby UPF 29
Configuring VPP Monitor on Active UPF and Standby UPF 30
Preventing User Plane Function Switchback 30
Preventing Dual Active Error Scenarios 31
Resetting Sx/N4 Monitor Failure 31
Changing UPF State from Pending-Active to Active 31
Monitoring and Troubleshooting 32
Show Command(s) and/or Outputs 32
show srp monitor bfd 32

show srp monitor bgp 32
show srp monitor sx 33
show srp monitor vpp 34
show srp statistics 34
L2 ICSR Support 35
Feature Description 35
SRP Loopback Interface 35
How it Works 36
State Machine 37
Configuring L2 ICSR in SRP Context 38
Configuring C-C and C-R BFD from Primary and Backup Chassis 38
Verifying L2 ICSR Configuration 38
OAM Support 40
Show Command Support 40

CHAPTER 4 **ADC for Dynamic Rules** 41
Feature Description 41
How it Works 42
Limitations 43
OAM Support 43
Show Command Output 43
show subscribers user-plane-only full all 43
show subscribers user-plane-only callid callid pdr full all 45

CHAPTER 5 **APN ACL Support** 49
Feature Summary and Revision History 49
Summary Data 49
Revision History 49
Feature Description 50
Rule(s) 50
Actions 50
Criteria 51
Rule Order 52
Limitations 52

Configuring ACL 52
Verifying ACL Configuration 53
IP Source Violation 53
Gating Control 53

CHAPTER 6 APN AMBR Traffic Policing 55
Feature Summary and Revision History 55
Summary Data 55
Revision History 55
Feature Description 56
Limitations 56
Configuring the APN AMBR Traffic Policing Feature 56
Monitoring and Troubleshooting 57
Show Commands and/or Outputs 57

CHAPTER 7 Bulk Statistics Support 59
Feature Summary and Revision History 59
Summary Data 59
Revision History 60
Feature Description 60
Bulk statistics for per-QCI drop counters 62

CHAPTER 8 Charging Support 63
Feature Summary and Revision History 63
Summary Data 63
Revision History 63
Feature Description 64
Offline Charging Events Reporting over N4 64
Online Charging Support over N4 64
How it Works 65
Call Flows 65
PFCP Session Establishment Procedure 65
PFCP Session Modification Procedure 66
PFCP Session Reporting Procedure 67

PFCP Session Deletion Procedure 67
IEs Supported for Offline Charging Reporting 68
IEs Supported for Online Charging Reporting 69
Usage Reporting in PFCP Modification Response 71
Usage Reporting for Online and Offline Charging 71
Usage Reporting with Rating-Group and Service ID 71
Implementing the QAURR Flag 72
Supported Functionality and Limitations 72
PTT no-quota Limited Pass 73
PTT quota exhaust Limited Pass 74
Tariff Time Support 75
TCP Maximum Segment Size 75
Configuring Credit Control for Usage Reporting 75
Configuring ACS Rulebase for Usage Reporting 76
Monitoring and Troubleshooting 79
Show Commands and/or Outputs 79
show-user-plane-service statistics rulebase name <name> 79

CHAPTER 9 **Cisco Ultra Traffic Optimization with VPP 81**
Feature Summary and Revision History 81
Summary Data 81
Revision History 81
Feature Description 82
RCM Support 82
Sending the GBR or MBR Values to Cisco Ultra Traffic Optimization 83
Cisco Ultra Traffic Optimization Library Deinitialization 83
How it Works 83
Architecture 83
Limitations 85
Show Commands and Outputs 85
Show Commands and Outputs 85
Bulkstats 89
Sample Configuration 91

CHAPTER 10 **Collection and Reporting of Usage Data over N4 Interface** 93

Feature Summary and Revision History 93

Summary Data 93

Revision History 93

Feature Description 94

How it Works 94

Standards Compliance 95

Configuration to Collect and Report Volume Measurement over N4 Interface 95

Configuring Charging Action for a Required Billing Action 95

Associating a Charging Action with a Rulebase 96

CHAPTER 11 **Control Plane-Initiated N4 Association Support** 97

Feature Summary and Revision History 97

Summary Data 97

Revision History 97

Feature Description 98

SMF initiated N4 Association Setup Procedure 98

How it Works 98

Call Flows 98

Session Management Function Initiated N4 Association Setup Procedure 98

Configuring the CP-Initiated N4 Association Setup Feature 98

CP-Initiated N4 Association Setup Feature OAM Support 99

Show Command Support 99

CHAPTER 12 **Converged Datapath** 101

Feature Summary and Revision History 101

Summary Data 101

Revision History 101

Feature Description 102

Architecture 102

How it Works 103

SxDemuxMgr 103

SessMgr 103

Datapath 103
Charging 103
Call Flows 104
Initial Attach with SGW-C/cnSGW and SMF/IWF 104
5G to 4G Handover with Collapsed UPF 106
Intra S-GW Handover with Collapsed UPF 108
Idle/Active DDN Handling with Collapsed UPF 109
IDFT Handling during S1 Handover 110
S-GW Relocation with Same SGW-U 111
WiFi to LTE Handover 113
Limitations 115
Charging for Converged Calls 116
Feature Description 116
How it Works 116
Limitations 117
Configuring eGCDR 117
Configuring GTPP Group 118
Generating Usage Report with Volume Count for Converged Sxa Calls 118
Configuration Example 118
Configuring Converged Datapath 118
Enabling Converged Datapath at UPF 118
Configuring Remote Peers for Sxa and N4 119
Configuring User Plane Service for Sxa and N4 119
Monitoring and Troubleshooting 119
Show Commands and/or Outputs 119
show subscribers user-plane-only all 119
show subscribers user-plane-only full all 120
show user-plane-service statistics all 120

CHAPTER 13 **Deep Packet Inspection and Inline Services** 121
Feature Summary and Revision History 121
Summary Data 121
Revision History 121
Feature Description 122

DNS Server Readdressing 122
How it Works 124
DSCP Marking for Downlink and Uplink Packets 124
Transport Level Marking IE 125
Transport Level Marking Options IE 125
Inner Packet Marking IE 125
Traffic Readdressing or Redirecting 126
Redirect Information IE 127
Supported Inline Services 127
Application Detection and Control 127
QUIC IETF Implementation 128
Configuring QUIC IETF 128
Statistics 128
SCONE support 129
Configure SCONE support 130
SCONE monitoring and troubleshooting 131
App-based Tethering Detection 132
Feature Description 132
Configuring App-based Tethering Detection 132
Monitoring and Troubleshooting the App-based Tethering Detection 133
Content Filtering 133
DNS Snooping 134
Event Data Records 135
Feature Description 136
How It Works 136
Configuring Event Data Records 139
Monitoring and Troubleshooting 142
Firewall Support 144
Feature Description 144
Subscriber Firewall Configuration 144
Monitoring and Troubleshooting 146
Flow Idle Timeout Randomization 148
Configuring Flow Idle Timeout Randomization in ACS 148
HTTP URL Filtering 149

IP Readdressing 152
Configuring IP Readdressing 153
Show Commands 153
L7 Protocol 154
DNS 154
FTP 154
HTTP 154
HTTPS 162
RTP/RTSP 162
SIP 162
Monitoring and Troubleshooting 162
NAT Support on N4 164
Feature Description 164
Configuring NAT 165
Monitoring and Troubleshooting 166
Tethering Detection 171
Feature Description 171
Configuring Tethering Support 171
Monitoring and Troubleshooting 172
RTP Dynamic Flow Detection 173
Rule-matching for Bearer-specific Filters 173
Enablement of IP Source Violation Detection for Uplink Packets 175
How Enabling of IP Source Violation for Uplink Packets Works 176
Verifying IP Source Violation for Uplink Packets 176
URL Blockedlisting 179
Feature Description 179
How it Works 179
Configuring URL Blockedlisting 180
Monitoring and Troubleshooting 181
Configuring the Static and Pre-Defined Rules 183
Configuring ACS Ruledef for L7 Protocols for DPI 184
Charging Action Configuration for L7 Protocols for DPI 186

CHAPTER 14 **Device ID in EDNS0 Records** 189

Feature Summary and Revision History 189
Summary Data 189
Revision History 189
Feature Description 190
How it Works 190
Process Flow 191
EDNS0 Packet Format 191
EDNS0 with IP Readdressing 192
Behavior and Restrictions 192
Limitation 193
Configuring EDNS Format and Trigger Action 193
Sample Configuration 195
Monitoring and Troubleshooting 195
Show Commands and Outputs 196
Bulk Statistics 197

CHAPTER 15 **Downlink Data Notification** 199
Feature Summary and Revision History 199
Summary Data 199
Revision History 199
Feature Description 200
How It Works 200
Downlink Data Notification – Delay (DDN-D) Support 200
5G SMF Calls 201
DDN Throttling Support 201
No User Connect Timer Support 202
DDN Call Flows 203
DDN Success Scenario 203
DDN Failure Scenario 204
No User Connect Timer Support 205
DDN Delay Timer 206
Sx Interface 207
Limitations 209
DDN Throttling for non-Release 10 Compliant MME 209

DDN Throttling for Release 10 Compliant MME 211
Idle Timer for SAE-GW Sessions 212
Limitations 212
Configuring Idle Timer for SAE-GW Sessions 212
S-GW Session Idle Timeout 213
Configuring Session Idle Timeout 213
Show Commands Input and/or Outputs 213
show subscribers user-plane-only full all 214
show subscribers user-plane-only full callid <call_id> 214

CHAPTER 16 DSCP Markings For Collapse Calls 215
Feature Summary and Revision History 215
Feature Description 216
DSCP Markings for 5G Calls 216
DSCP Markings for 4G Collapsed Datapath Calls 216
How It Works 216
SessMgr SMF Changes 217
Configuration 219
Monitoring and Troubleshooting 220
Show Commands Outputs 220

CHAPTER 17 Dynamic and Static PCC Rules 221
Feature Summary and Revision History 221
Summary Data 221
Revision History 222
Feature Description 222
How it Works 222
Predefined PCC Rules Support 222
Provisioning of Predefined PCC Rules 222
Dynamic PCC Rules Support 223
Policing 224
Bandwidth Policy Configuration Limits 226
Rate Limiting for Static and Predefined Rules 226
Rate Limiting for Dynamic Rules 227

Standards Compliance 228
Configuring the URR IDs 228
Threshold Configuration 229

CHAPTER 18 ECS Regular Expression 231
Feature Summary and Revision History 231
Summary Data 231
Revision History 231
Feature Description 232
How It Works 232
Configuring Regex Rule 233
Sample Configuration 234
Monitoring and Troubleshooting 234
Show Commands and Outputs 234

CHAPTER 19 GTP-U Support 235
Feature Summary and Revision History 235
Summary Data 235
Revision History 235
Feature Description 236
How it Works 237
Call Flows 237
Initial Attach on E-UTRAN via MME and S-GW 237
5G to EPS Handover with N26 Interface 238
Error Indication Handling on UPF 239
GTP-U Path Failure Support at UPF 239
Enabling Dual Stack Support on S5u/N3 Interface 240
Disabling UDP Checksum 240
Disabling UDP Checksum 240

CHAPTER 20 Heartbeat Support for N4/Sx Interface 241
Feature Summary and Revision History 241
Summary Data 241
Revision History 241

Feature Description 242
How It Works 242
Path Failure Detection 242
Path Failure Handling 243
Configuring Heartbeat for N4/Sx Interface 243
Enabling Heartbeat for Sx Interface 243
Configuring Detection Policy for Path Failure 244
Monitoring and Troubleshooting 244
Show Command(s) and/or Outputs 244
show sx-service all 244
show sx-service statistics all 245
Disconnect Reasons 245
SNMP Traps 245

CHAPTER 21 **Idle Mode Buffering and Paging** 247
Feature Summary and Revision History 247
Summary Data 247
Revision History 247
Feature Description 248
How it Works 248
Provisioning of Buffering Action Rule in the UPF 248
Buffering Action Rule Call Flow 248
Downlink Data Report for First DL Packet 249
Paging Policy Differentiation 249
Paging Policy Indicator (PPI) 249
Frame Format for the PDU Session User Plane Protocol 250
QoS Flow Identifier (QFI) 250
Paging Policy Presence 250
Paging Policy Indicator 250

CHAPTER 22 **Indirect Forwarding Tunnel** 251
Revision History 251
Feature Description 251
How It Works 251

Call Flow 251
Supported Functionality 254
Limitations 254
Configuring Indirect Forwarding Tunnel 254
Enabling Indirect Forwarding Tunnel Feature 255
Verifying the Indirect Forwarding Tunnel Feature 255
show sgw-service name <service_name> 255
Monitoring and Troubleshooting 255
Show Commands Input and/or Outputs 255
show subscribers saegw-only full all 255
show subscribers user-plane-only callid <call_id> pdr all 255
show subscribers user-plane-only full all 256

CHAPTER 23 IP Allocation through DHCP Server 257
How UPF allocates the UE IP through DHCP server 258
DHCP subnet selection criteria 260
Criteria for DHCP configuration validation 261
Validation of IPv4 and IPv6 addresses 261
Enabling DHCP-based IP allocation 263
Start the DHCPv4 or DHCPv6 service 263
Configure slot/port to enable discovery of DHCP server 264
Limitations 267

CHAPTER 24 IPsec Support for IPv6 269
Feature Summary and Revision History 269
Summary Data 269
Revision History 269
Feature Description 270
IPsec AH and ESP 270
IPsec Transport and Tunnel Mode 270
IPsec Terminology 270
Crypto Access Control List 270
Transform Set 271
ISAKMP Policy 271

Crypto Map 271
Crypto Template 271
Supported Algorithms 271
Limitations and Restrictions 272
Example Configurations 272
Monitoring and Troubleshooting 273
Show Commands 273

CHAPTER 25 IPv6 Prefix Delegation 275
IPv6 Prefix Delegation 276
Licensing 276
Guidelines and Limitations of IPv6 Prefix Delegation 277
Delegation of IPv6 Prefix 277
IPv6 Prefix Release 279
How IPv6 Prefix Delegation Works 279
Session Modificaiton Request based Prefix Allocaiton 280
Configure DHCP Service Name for Prefix Delegation 281
Verify IPv6 Prefix Delegation 283
show subscribers user-plane-only full all callid 283
show dhcpv6 statistics 283

CHAPTER 26 LTE - Wi-Fi Seamless Handover 285
Feature Summary and Revision History 285
Summary Data 285
Revision History 285
Feature Description 285
How It Works 286
EPC to Non-3GPP Untrusted Wi-Fi Handover Call Flow 286
Non-3GPP Untrusted Wi-Fi to EPC Handover Call Flow 289
Non-3GPP Interworking Function 292
ICSR and Session Recovery 293
Limitations 293
Standards Compliance 293

CHAPTER 27 **Monitor Subscriber** 295

Feature Summary and Revision History 295

Summary Data 295

Revision History 295

Feature Description 296

How It Works 297

SMF and UPF Interactions 297

Subscriber Tracing from SMF 298

Fastpath and Slowpath Packet Captures 302

UPF SessMgr Functionality 302

Multi PDN Multi-trace 303

MonSub Statistics 303

X-Header 303

Configuration Procedure for Monitor Subscriber on UPF 303

Monsub CLI Options 304

Context, CDRMOD, and Hexdump Interaction for Monitor Subscriber 306

PCAP File Name Convention 306

PCAP File Location 309

Limitations 309

Configuring the Hexdump Module for MonSub in UPF 311

Configuring MonSub Poll Timer 311

Configuring MonSub File Name 311

Monitoring and Troubleshooting 312

SNMP Traps 312

CHAPTER 28 **MPLS Support on UPF** 315

Feature Summary and Revision History 315

Summary Data 315

Revision History 315

Feature Description 316

How it Works 316

MPLS-CE Connected to PE 317

Sample Configuration 318

VPN-related CLI Commands 320
Monitoring and Troubleshooting 323
Show Commands and Outputs 323
show mpls ftn vpp 323

CHAPTER 29 **Multiple cnSGW Support** 325
Feature Summary and Revision History 325
Summary Data 325
Revision History 325
Feature Description 326
How it Works 326
Configuring Multiple SMF/cnSGWs 326
Configuring Multiple SMF/cnSGWs on UPF 327
Monitoring and Troubleshooting 327
Show Commands and/or Outputs 327
show subscribers user-plane-only full all 327
show sx peers 327

CHAPTER 30 **Multiple N4/Sx Interfaces** 329
Feature Summary and Revision History 329
Summary Data 329
Revision History 329
Feature Description 330
How it Works 330
Configuring Multiple N4/Sx Interfaces 331
Configuring Multiple Peer Nodes 331
Monitoring and Troubleshooting 332
Show Commands and/or Outputs 332
show ip chunks 332
show ipv6 chunks 332
show subscribers user-plane-only full all 332
show sx peers 332
show user-plane-service statistics peer-address <address> 332

CHAPTER 31 **Nexthop Forwarding Support** 333
Revision History 333
Feature Description 333
How It Works 333
Architecture 333
Limitations 338
Configuring Next Hop Forwarding Support 338
Configuring Next Hop Forwarding Through Charging Action 338
Configuring Next Hop Forwarding Through DNN Profile 339
Configuring Next Hop Forwarding at IP Pool Through IPAM Profile for IPv4 Addresses 339
Configuring Next Hop Forwarding at IP Pool Through IPAM Profile for IPv6 Addresses 340
Configuration Example 341
VLAN Segmentation Using Next Hop Forwarding 342
Feature Description 342
How it Works 343
Enabling VLAN Segmentation Using Next Hop Forwarding 343
Configuration Example 343
Configuring Network Reachability Server on UPF 344
Configuration Example 345
OAM Support 345
Monitoring Support 345
Monitoring and Troubleshooting 346
Show Commands and Outputs 347

CHAPTER 32 **N:M Redundancy and Redundancy Configuration Manager** 349
Feature Summary and Revision History 349
Summary Data 349
Revision History 349
Feature Description 350

CHAPTER 33 **N3 Transfer of PDU Session Information** 351
Feature Summary and Revision History 351
Summary Data 351

Revision History 351
Feature Description 351
How it Works 352
Transfer of PDU Session Information for Downlink Data Packets 352
Transfer of PDU Session Information for Uplink Data Packets 352
PDU Session Information Frame IEs 353
Standards Compliance 354
Limitations 354

CHAPTER 34 **N4 Interface Compliance with 3GPP Specification** 355
Feature Summary and Revision History 355
Summary Data 355
Revision History 355
Feature Description 356
Averaging Window 356
Paging Policy Indicator 356
Outer Header Creation 357
Outer Header Removal 358

CHAPTER 35 **N4 Interface Configuration** 361
Feature Summary and Revision History 361
Summary Data 361
Revision History 361
Feature Description 362
Uniformity in compression at N4 and Sx interfaces 362
Configuring N4 Interface 363
Identifying an N4 Interface 363
Modification of N4-type Parameters in an Sx Service 364
Statistics 364
show control-plane-group 364
show sx-service all 364
show subscribers user-plane-only all 364
show user-plane-service statistics all 365
show subscribers user-plane-only seid number pdr all 365

show subscribers user-plane-only callid number pdr full all 365

CHAPTER 36 N4/Sx over IPSec 367

Feature Summary and Revision History 367

Summary Data 367

Revision History 367

Feature Description 368

Limitations 370

Recommended Timers 370

Recommended Configurations 370

Example Configurations in SMF 371

Example Router Configurations 374

Example Configurations in UPF 375

Example SRP Configurations 376

Sample Configurations 377

Monitoring and Troubleshooting 381

CHAPTER 37 N4 Session Management, Node Level, and Reporting Procedures 385

Feature Summary and Revision History 385

Summary Data 385

Revision History 385

Feature Description 386

N4 Session Management, Node Level, and Reporting Procedures 386

N4 Node-level Procedures 386

N4 Session Management 386

N4 Session/Node-level Reporting Procedures 386

Relationships 387

End Marker Support 387

UEs IPv4, IPv6, and IPv4v6 Support 387

How it Works 387

N4 Node-level Procedure Call Flows 387

N4 Association Setup Procedure Call Flow 387

N4 Association Update Procedure Call Flow 388

N4 Association Release Procedure Call Flow 389

N4 Heartbeat Procedure 389
N4 Session Management Procedures Call Flows 389
N4 Session Establishment Call Flow 389
N4 Session Modification Call Flow 390
N4 Session Delete Call Flow 391
N4 Session/Node Level Reporting Procedure Call Flows 391
Session Level Reporting Due to the GTP-u Error Indication Call Flow 391
Node-level Reporting Procedure due to GTP-u Path Failure Call Flow 392
PDN Update Procedure - eNodeB F-TEIDu 394
UEs IPv4, IPv6, and IPv4v6 Support Call Flows 395
N4 Session Establishment and Modification Procedure for IPv6 Call Flow 395
N4 Session Establishment and Modification Procedure for IPv4v6 Call Flow 396
Configuring the N4 Session/Node Level Reporting Procedures 397
Enabling the GTP-u Echo Request Procedure 397
Verifying the N4 Session/Node Level Reporting Procedure Configuration 398
N4 Session Node Level Reporting Procedure OA and M Support 398
SNMP Traps 400
Bulk Statistics 400
Enhanced PFCP Association Release Procedure for Graceful Session Termination 401
UPF-initiated PFCP Session Release 402
Clear the Subscriber Session in UPF 402
Recovery Handling in UPF 403
Verify the Cause Level Statistics 403
UPF-initiated Enhanced PFCP Association Release 404
Enable the EPFAR Feature 404
Initiate the Association Release for Sx Service 405
Enable the Association Release Timer 405

CHAPTER 38 **NRF Support** 407
Feature Summary and Revision History 407
Summary Data 407
Revision History 407
Feature Description 408
NRF Management Services 408

How it Works 409
UPF Registration 409
UPF Heartbeat 409
UPF DeRegistration 409
Standards Compliance 409
Configuring NRF Management Services 409
Monitoring and Troubleshooting 411
Show Commands and/or Outputs 411

CHAPTER 39 **Password Expiration Notification 413**
Feature Summary and Revision History 413
Summary Data 413
Revision History 413
Feature Description 414
Upgrading and Downgrading Procedures using Save Configuration Command 415

CHAPTER 40 **QCI 80 Support on UPF 417**
Feature Summary and Revision History 417
Summary Data 417
Revision History 417
Feature Description 418
How it Works 418
Dynamic QoS Flow Establishment based on Detected Traffic 418
Call Flow 418
Limitations 420
Configuring ADC Rule 420
Monitoring and Troubleshooting 421

CHAPTER 41 **QoS Group of Ruledefs Support 423**
Feature Summary and Revision History 423
Summary Data 423
Revision History 423
Feature Description 424
How it Works 424

Limitations 426
OAM Support 426
Show Commands Support 426
show subscribers user-plane-only callid call_id far full all 426
show subscribers user-plane-only callid call_id qer full all 427
show subscribers user-plane-only callid call_id qos-group all 427
show subscribers user-plane-only callid call_id qos-group statistics 427
show subscribers user-plane-only full all 428
show user-plane-service qos-group-of-ruledefs 429
show user-plane-service statistics qos-group sessmgr 429

CHAPTER 42 **Roaming Support** 431
Feature Summary and Revision History 431
Summary Data 431
Revision History 431
Feature Description 432
Relationships 433
Architecture 433
Inter-PLMN Mobility Support 434
How it Works 435
Standards Compliance 451
Limitations 451
Configuring the HR Roaming Support for UPF 452
Roaming Status in Sub-Params IE 453
Source Interface Type IE 454
Monitoring and Troubleshooting 454
show user-plane-service statistics all 454
show subscribers user-plane-only full all 455
show subscribers user-plane-only callid pdr full all 455
show user-plane-service gtpu statistics gtpu-service n9-egress 456
show user-plane-service gtpu statistics gtpu-service n9-ingress 458

CHAPTER 43 **Secure Group Tag-based Access Control** 461
Feature Summary and Revision History 461

Summary Data 461
Revision History 461
Feature Description 462
Cisco ISE 463
Security Group Tag 463
Security Group Tag-based Access Control List 464
Security Group Tag Matrix 464
Fastpath Stream 464
How it Works 464
Viewing Fastpath Stream Details 464
User Access Control through SGACL 465
How SGACL for ISE Integration Works 467
Call Flow 467
Limitations 470
Secondary ISE server integration on UPF 471
How secondary ISE server integration works 473
Limitations of secondary ISE server integration 474
SGACL Configuration for ISE Integration 474
Enable API Manager 474
Configure ISE Server profile 475
Manually switch to active ISE server 477
Associate ISE Server Profile 477
Refresh D-SGT Column 478
Monitoring and Troubleshooting 478
Verify SGACL with SGT Integration 478
show subscribers user-plane-only full callid callid_value 479
show subscribers user-plane-only callid callid_value flows full 480
show user-plane-service sgt-column summary 480
show user-plane-service sgt-column dsgt 480
show user-plane-service sgacl name 481
show user-plane-service statistics sgacl all 481
show user-plane-service statistics drop-counter 481
show apimgr statistics ise-server 481
SNMP trap 482

OAM Support 482

CHAPTER 44 Session Recovery 483

Feature Summary and Revision History 483

Summary Data 483

Revision History 483

Feature Description 483

How it Works 484

Configuring the System to Support Session Recovery 484

Enabling Session Recovery 484

Enabling Session Recovery on an Out-of-Service System 484

Enabling Session Recovery on an In-Service System 485

Disabling the Session Recovery Feature 486

Viewing Session Recovery Status 486

Viewing Recreated Session Information 487

CHAPTER 45 Session Report Rejection Procedure 489

Feature Summary and Revision History 489

Summary Data 489

Revision History 489

Feature Description 489

Relationships to Other Features 490

Call Flow 490

OAM Support 491

Show Command(s) and/or Output(s) 491

CHAPTER 46 Smart Licensing 493

Feature Summary and Revision History 493

Summary Data 493

Revision History 493

Overview 494

Cisco Smart Software Manager 495

Smart Accounts/Virtual Accounts 495

Smart Licensing Mode 495

Request a Cisco Smart Account 495
Software Tags and Entitlement Tags 496
Configuring Smart Licensing 499
Monitoring and Troubleshooting Smart Licensing 500

CHAPTER 47 **Software Management Operations** 501
Feature Summary and Revision History 501
Summary Data 501
Revision History 501
Overview 502
SNMP Traps 503
Limitations 503
Health Checks 503
Build Upgrade 505
UPF Upgrade 507
UPF Downgrade 507

CHAPTER 48 **Standard QCI Support** 509
Feature Summary and Revision History 509
Summary Data 509
Revision History 509
Feature Description 509
Limitations 510
Configurations 510

CHAPTER 49 **System Logs** 511
Feature Summary and Revision History 511
Summary Data 511
Revision History 512
System Log Types 512
Configuring Event Logging Parameters 513
Configuring Event Log Filters 513
Exec Mode Filtering 513
Global Configuration Mode Filtering 515

Configuring Syslog Servers 516
Configuring Active Logs 518
Specifying Facilities 519
Configuring Trace Logging 528
Configuring Monitor Logs 528
Enabling Monitor Logs 528
Disabling Monitor Logs 528
Viewing Logging Configuration and Statistics 529
Viewing Event Logs Using the CLI 529
Configuring and Viewing Crash Logs 530
Crash Logging Architecture 531
Configuring Software Crash Log Destinations 531
Viewing Abridged Crash Log Information Using the CLI 531
Reducing Excessive Event Logging 532
Configuring Log Source Thresholds 533
Checkpointing Logs 534
Saving Log Files 534
Event ID Overview 536
Event Severities 544
Understanding Event ID Information in Logged Output 544

CHAPTER 50 **UPF Ingress Interfaces** 547
Feature Summary and Revision History 547
Summary Data 547
Revision History 547
Feature Description 548
UPF Ingress Interface Type Configuration 548
Configuring UPF Ingress Interface Type 548
Configuration Example 549
Configuring Echo Response for N9 Interface 549
Separating N3 and N9 BGP Routes 549
Verifying the UPF Ingress Interface Type Configuration 553
show user-plane-service all 553
Multiple GTPU IP Addresses Support 554

Feature Description 554
How it Works 554
Limitations 555
Configuring Multiple Bind Addresses on N3 Interface 555
Configuration Example 556
Configuration Verification 556
OAM Support 557
Monitoring Support 557

CHAPTER 51 UPF Local Configuration 561
Feature Summary and Revision History 561
Summary Data 561
Revision History 561
Feature Description 562
How it Works 562
Configuring the Local Configuration Support for UPF 563
TACACS+ TLS support 564
Configure TACACS+ TLS 564
TACACS+ TLS monitoring and troubleshooting 565

CHAPTER 52 Load and Overload Control Over N4/Sx Interface 567
Feature Summary and Revision History 567
Summary Data 567
Revision History 568
Feature Description 568
How Load and Overload Control Works 569
Load and Overload Factor Calculation 569
Load and Overload Reporting to CP-F 570
Load Information Reporting to CP-F 570
Overload Information Reporting to CP-F 571
Session Clear in Self-Protection Mode 571
Enable Load and Overload Control on N4 or Sx Interface 572
Configure the Load Control Profile 572
Configure the Overload Control Profile 573

Associate the Load and Overload Control Profiles with User Plane Service 574
Monitoring and Troubleshooting 575
Show Commands and Outputs 575
show upf-load-control-profile 575
show upf-overload-control-profile 575
show user-plane-service statistics name 576
show user-plane-service statistics all 577
show user-plane-service name 578
show user-plane-service all 579
show sx-service statistics all 579
Bulkstats 579
SNMP Traps 581

CHAPTER 53 UPF Usage Monitoring over PCF 583
Feature Summary and Revision History 583
Summary Data 583
Revision History 583
Feature Description 584
Usage Reporting 584

CHAPTER 54 Variable Relinquishing of CPU Cycles 585
VRCC Overview 585
Benefits of VRCC 586
How VRCC Works 586
Configure VRCC 587

CHAPTER 55 Virtual Routing and Forwarding 589
Feature Summary and Revision History 589
Revision History 589
Feature Description 589
Overlapping IP Pool 590
VRF Name as Identifier 591
Limitations and Restrictions 591
Configuring VRF 592

VLAN Segmentation Using VRF 595
Feature Description 595
Enabling VLAN Segmentation Using VRF 595
Configuration Example 596
Monitoring and Troubleshooting 596
Show Commands and Outputs 596
show ip chunks 596
show ipv6 chunks 597
show ip bgp vpnv4 597
show ip bgp vpnv6 598
show mpls ilm 598

CHAPTER 56 **Voice over New Radio** 599
Feature Summary and Revision History 599
Summary Data 599
Revision History 599
Feature Description 600
How it Works 600
VoNR Call Flow for UPF 600

CHAPTER 57 **WPS Prioritization** 603
Feature Description 603
How it Works 604
Limitations 606
Message Prioritization for Overload Scenarios 607
Configuring WPS Prioritization 607
Enabling WPS Prioritization 607
Mapping the Session Priority Profile 608
OAM Support 609
Show Commands Support 609
show subscribers user-plane-only full all 609
show user-plane-service all 609
show user-plane-service session-priority-profile name | all 609
show user-plane-service statistics all 610

show user-plane-service statistics session-priority-profile 610

CHAPTER 58 X-Header Insertion and Encryption 613

Feature Summary and Revision History 613

Summary Data 613

Revision History 613

Feature Description 614

How it Works 614

X-Header Insertion 614

X-Header Encryption 614

Limitations 615

Configuring X-Header Insertion and Encryption 615

Configuring X-Header Insertion 615

Configuring X-Header Encryption 616

Verifying X-header Insertion and Encryption Configuration 617

Monitoring and Troubleshooting X-Header Insertion and Encryption Configuration 617

Anti-Spoofing Support 618

Feature Description 618

Limitations 618

Enabling Anti-Spoofing in X-Header 618

OAM Support 619

CHAPTER 59 UPF Troubleshooting Information 621

Debug Logging 621

Monitoring CLI 622

Monitoring Protocol 623

RAT Type-based Statistics 623

Subscriber Level CLI 627

show subscribers user-plane-only full msid msid_value 628

show subscribers user-plane-only callid callid-value rule-match-info { [full] all | rulename rulename} 633

VPP Statistics 635

show subscribers user-plane-only callid callid_value drop-statistics 636

show user-plane-service statistics drop-counter 637

show subscribers user-plane-only callid callid_value flows full 638
show subscribers user-plane-only callid callid_value flows flow-id flow_id 638
SNMP Support 639
Troubleshooting UPF Features 640
Configure Stats Profile for QCI Drop Statistics 642
NSO configuration considerations 643
Per-QCI Drop statistics commands 643

CHAPTER 60 **Sample UPF Configuration** 647
Sample Configuration 647



About this Guide

Note Control and User Plane Separation (CUPS) represents a significant architectural change in the way StarOS-based products are deployed in the 3G, 4G, and 5G networks. This document provides information on the features and functionality specifically supported by this 5G UPF product deployed in a 5G network. It should not be assumed that features and functionality that have been previously supported in legacy or non-CUPS products are supported by this product. References to any legacy or non-CUPS products or features are for informational purposes only. Furthermore, it should not be assumed that any constructs (including, but not limited to, commands, statistics, attributes, MIB objects, alarms, logs, services) referenced in this document imply functional parity with legacy or non-CUPS products. Please contact your Cisco Account or Support representative for any questions about parity between this product and any legacy or non-CUPS products.

Note The documentation set for this product strives to use bias-free language. For purposes of this documentation set, bias-free is defined as language that does not imply discrimination based on age, disability, gender, racial identity, ethnic identity, sexual orientation, socioeconomic status, and intersectionality. Exceptions may be present in the documentation due to language that is hardcoded in the user interfaces of the product software, language used based on RFP documentation, or language that is used by a referenced third-party product.

This preface describes the *5G User Plane Function Guide*, how it is organized and its document conventions.

This guide describes the Cisco User Plane Function (UPF) and includes infrastructure and interfaces, feature descriptions, specification compliance, session flows, configuration instructions, and CLI commands for monitoring and troubleshooting the system.

• Conventions Used, on page xxxv

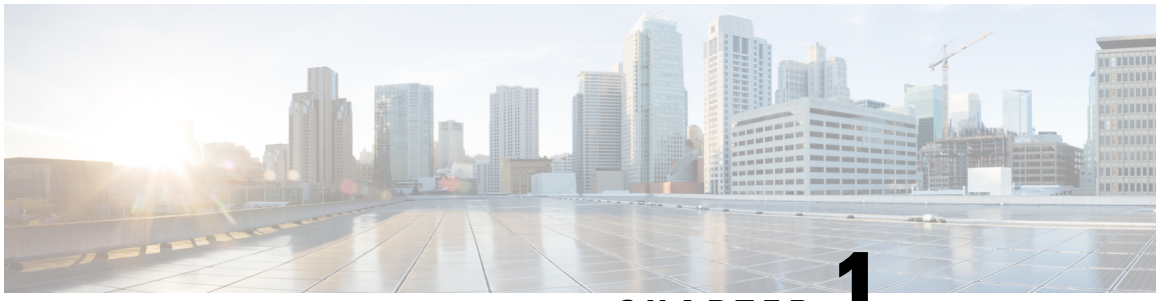
Conventions Used

The following tables describe the conventions used throughout this documentation.

Notice Type	Description
Information Note	Provides information about important features or instructions.

Notice Type	Description
Caution	Alerts you of potential damage to a program, device, or system.
Warning	Alerts you of potential personal injury or fatality. May also alert you of potential electrical hazards.

Typeface Conventions	Description
Text represented as a `screen display`	This typeface represents displays that appear on your terminal screen, for example: `Login:`
Text represented as **commands**	This typeface represents commands that you enter, for example: **show ip access-list** This document always gives the full form of a command in lowercase letters. Commands are not case sensitive.
Text represented as a **command** *variable*	This typeface represents a variable that is part of a command, for example: **show card** *slot_number* *slot_number* is a variable representing the desired chassis slot number.
Text represented as menu or sub-menu names	This typeface represents menus and sub-menus that you access within a software application, for example: Click the **File** menu, then click **New**



CHAPTER 1

5G Architecture

- Feature Summary and Revision History, on page 1
- Architectural Overview, on page 2
- Subscriber Microservices Infrastructure Architecture, on page 3
- Control Plane Network Function Architecture, on page 4

Feature Summary and Revision History

Summary Data

Table 1: Summary Data

Applicable Product(s) or Functional Area	• PCF • SMF • UPF
Applicable Platform(s)	SMI
Feature Default Setting	Not Applicable
Related Documentation	Not Applicable

Revision History

Table 2: Revision History

Revision Details	Release
First introduced.	Pre-2020.02.0

Architectural Overview

The Ultra Cloud Core is Cisco's solution supporting 3GPP's standards for 5G new radio (NR) standalone (SA) mode. These standards define various network functions (NFs) based on the separation of control plane (CP) and user plane (UP), such as CUPS. This separation enhances network performance and capabilities.

Control Plane Network Functions

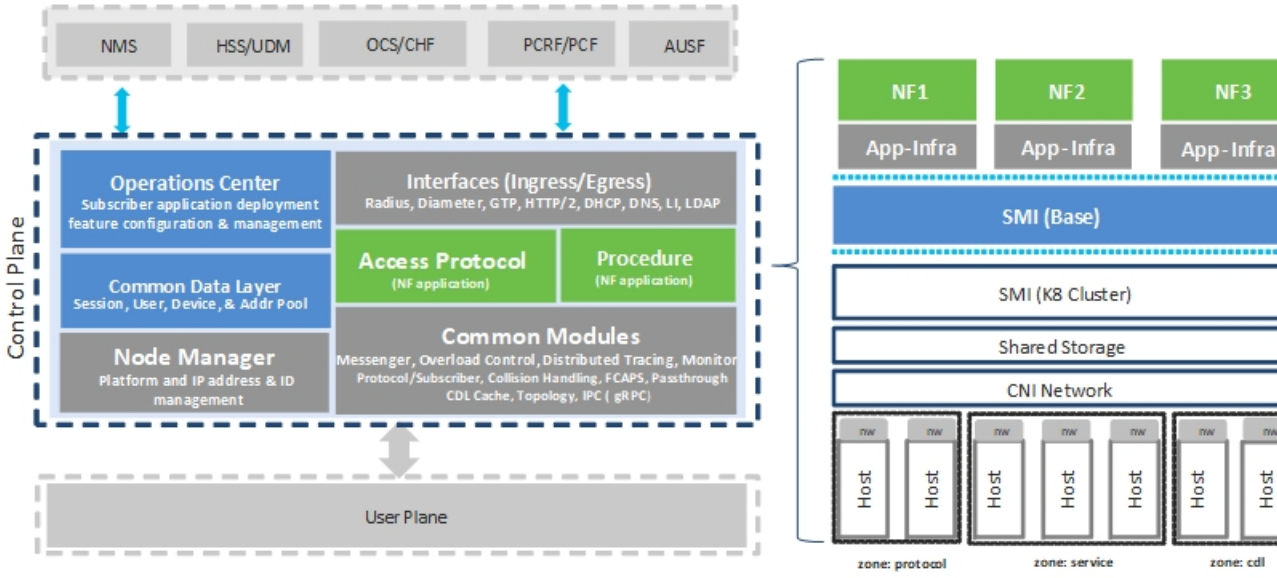
The Ultra Cloud Core's CP-related NFs use a common architecture with the following tenets:

- Cloud-scale—Fully virtualized for simplicity, speed, and flexibility.
- Automation and orchestration—Optimized operations, service creation, and infrastructure.
- Security—Multiple layers of security across the deployment stack from the infrastructure through the NF applications.
- API exposure—Open and extensive for greater visibility, control, and service enablement.
- Access agnostic—Support for heterogeneous network types (for example 5G, 4G, 3G, Wi-Fi, and so on).

Control plane NFs are designed as containerized applications, such as microservices. They deploy through the Subscriber Microservices Infrastructure (SMI).

The SMI defines the common application layers for functional aspects of the NF such as life-cycle management (LCM), operations and management (OAM), and packaging.

Figure 1: Ultra Cloud Core CP Architectural Components



User Plane Network Function

The 5G UP NF within the Ultra Cloud Core is the User Plane Function (UPF). Unlike the CP-related NFs, The 5G UPF uses the same Vector Packet Processing (VPP) technology as the Cisco 4G CUPS architecture. This commonality delivers a consistent set of capabilities between 4G and 5G, such as:

- Ultrafast packet forwarding.
- Extensive integrated IP Services such as Subscriber Firewall, Tethering, Deep-Packet Inspection (DPI), Internet Content Adaption Protocol (ICAP), Application Detection and Control (ADC), and header enrichment (HE).
- Integrated third-party applications for traffic and TCP optimization.

Subscriber Microservices Infrastructure Architecture

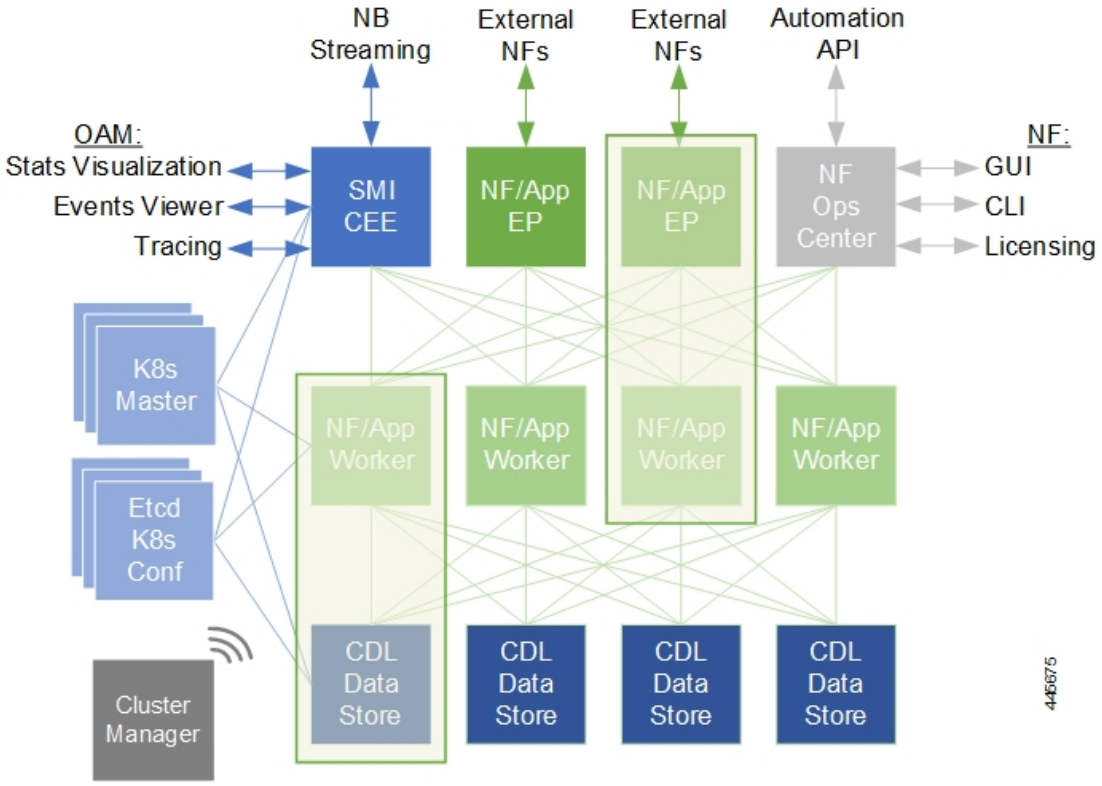
The Ultra Cloud Core (UCC) Subscriber Microservices Infrastructure (SMI) is a layered stack of cloud technologies. It enables rapid deployment and seamless life-cycle operations for microservices-based applications:

The SMI stack consists of the following components:

- SMI Cluster Manager—Creates the Kubernetes (K8s) cluster, creates the software repository, and provides ongoing LCM for the cluster, including deployment, upgrades, and expansion.
- Kubernetes Management—Includes the K8s primary and etcd functions, which provide LCM for the NF applications that are deployed in the cluster. This component also provides cluster health monitoring and resources scheduling.
- Common Execution Environment (CEE)—Offers utilities and OAM functionalities for Cisco Cloud native NFs and applications. These include licensing and entitlement functions, configuration management, and telemetry. It also provides alarm visualization, logging management, and troubleshooting utilities. Also, it provides consistent interaction and experience for all customer touch points and integration points in relation to these tools and deployed applications.
- Common Data Layer (CDL)—Provides a high performance, low latency, stateful data store, designed specifically for 5G and subscriber applications. This next generation data store offers high availability in local or geo-redundant deployments.
- Service Mesh—Provides sophisticated message routing between application containers, enabling managed interconnectivity, extra security, and the ability to deploy new code and new configurations in low risk manner.
- NB Streaming—Provides Northbound Data Streaming service for billing and charging systems.
- NF or Application Worker Nodes—The containers that comprise an NF application pod.
- NF or Application Endpoints (EPs)—The NFs or applications and their interfaces to other entities on the network
- Application Programming Interfaces (APIs)—Provides various APIs for deployment, configuration, and management automation.

The following figure depicts how these components interconnect to comprise a microservice-based NF or application.

Figure 2: SMI Components



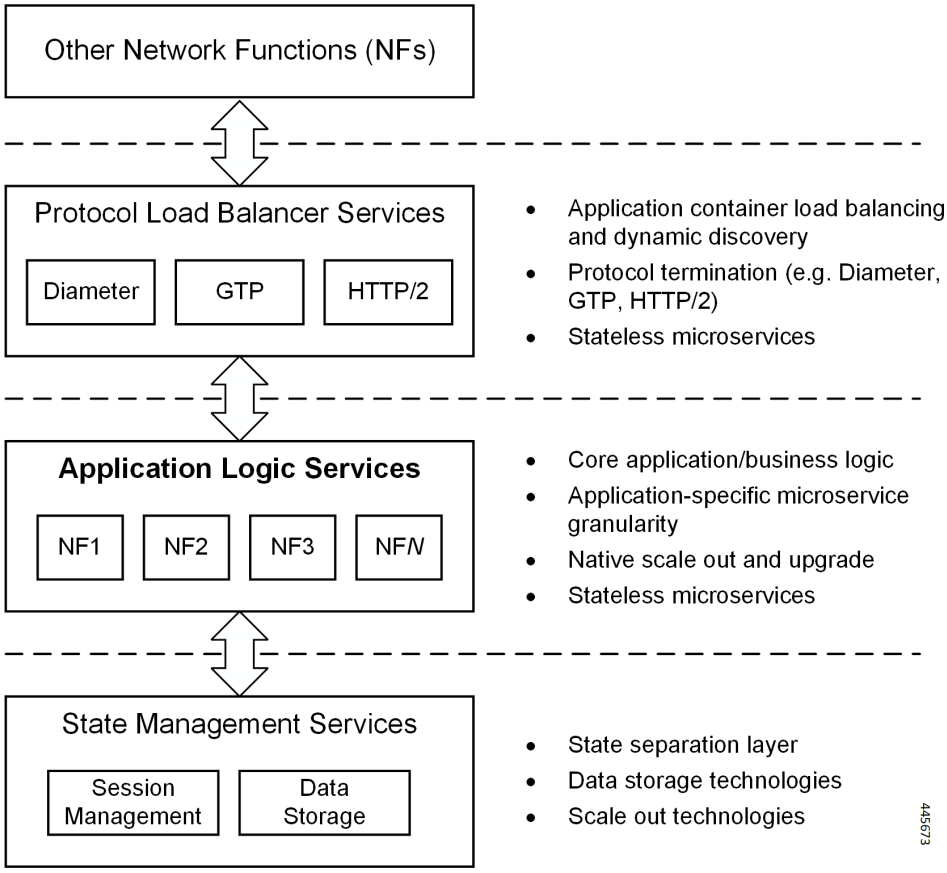
For more information on SMI components, see Ultra Cloud Core Subscriber Microservices Infrastructure and the related-documentation at *Deployment Guide* > *Overview* chapter.

Control Plane Network Function Architecture

Control plane (CP) NFs use a three-tiered architecture. This architecture leverages stateful or stateless capabilities in cloud native environments. The architectural tiers are as follows:

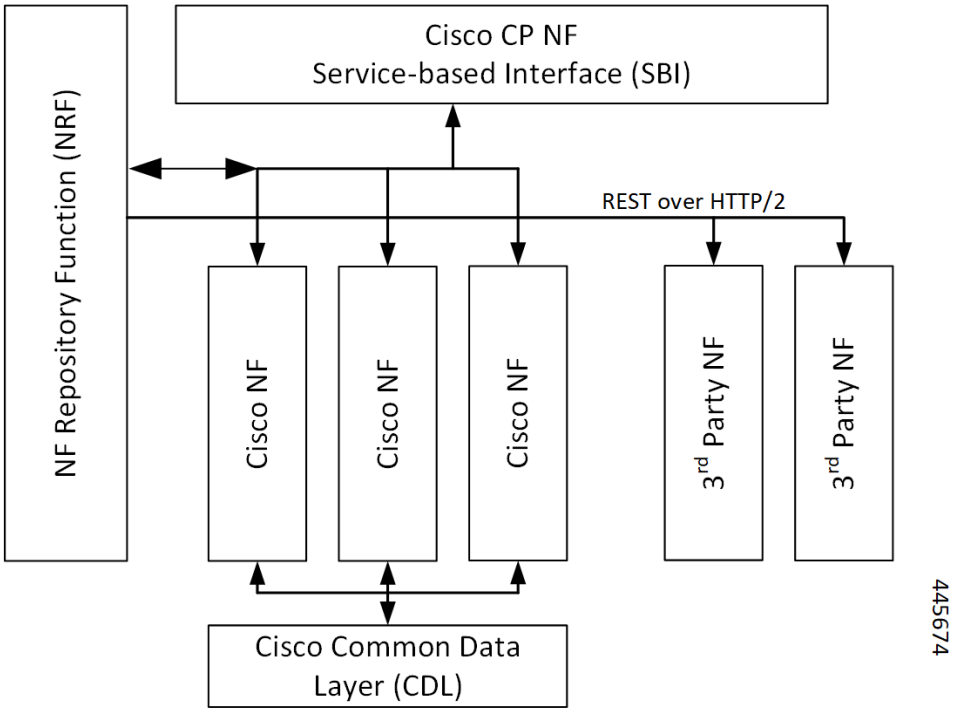
- Protocol Load Balancer Services—These are stateless microservices that are primarily responsible for dynamic discovery of application containers as well as for protocol proxy and termination. These include traditional 3GPP protocols and new protocols that are introduced with 5G.
- Applications Services—Responsible for implementing the core application or business logic, these are the stateless services that render the actual application based on the received information. This layer may contain varying degrees of microservice granularity. Application services are stateless.
- State management services—Enable stateless application services by providing a common data layer (CDL) to store or cache state information (for example session and subscriber data). This layer supports various data storage technologies from in-memory caches to full-fledge databases.

Figure 3: Control Plan Network Function Tiered Architecture

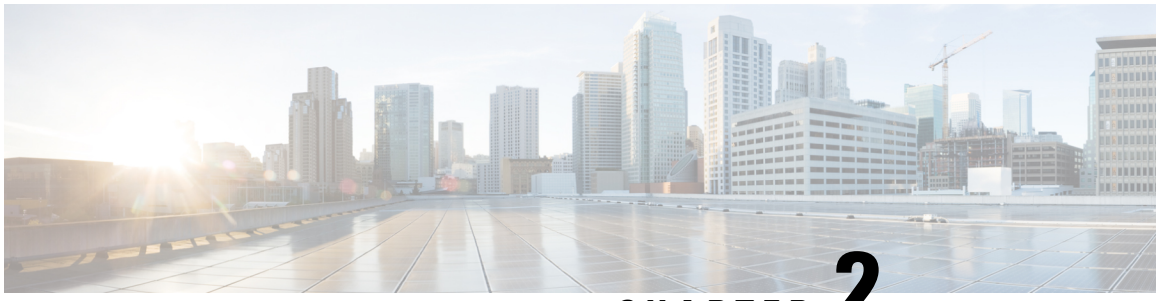


The three-tiered architecture on which Cisco CP NFs are designed fully support the 5G core (5GC) Service-based Architecture (SBA) defined by 3GPP. These NFs communicate using the Service-based Interface (SBI). They use HTTP/2 over TCP, as defined by 3GPP. Alternatively, Cisco's CP NFs can be configured to use gRPC over HTTP/2 for enhanced API performance between Cisco's NFs.

Figure 4: Cisco CP NF Service-based Architecture Support



For more information on the Cisco network functions, see their corresponding network function documentation.



CHAPTER 2

5G-UPF Overview

This chapter covers the following topics:

- Feature Summary and Revision History, on page 7
- Product Description, on page 8
- Use Cases and Features, on page 8
- Deployment Architecture and Interfaces, on page 12
- Standards Compliance, on page 17
- Limitations, on page 18

Feature Summary and Revision History

Summary Data

Table 3: Summary Data

Applicable Product(s) or Functional Area	5G-UPF
Applicable Platform(s)	VPC-SI SMI
Feature Default Setting	Disabled – License Required
Related Changes in this Release	Not Applicable
Related Documentation	Not Applicable

Revision History

Table 4: Revision History

Revision Details	Release
Added support for UPF cloud-native deployment.	2023.01.0
First introduced.	2020.02.0

Product Description

The User Plane Function (UPF) is one of the network functions (NFs) of the 5G core network (5GC). The UPF is responsible for packet routing and forwarding, packet inspection, QoS handling, and external PDU session for interconnecting Data Network (DN), in the 5G architecture.

UPF is a distinct Virtual Network Function (VNF) that offers a high-performance forwarding engine for the user traffic. Using Vector Packet Processing (VPP) technology, the UPF achieves ultra-fast packet forwarding while retaining compatibility with all the user plane functionality. For instance, Shallow Packet Inspection(SPI)/Deep Packet Inspection (DPI), traffic optimization, and inline services (NAT, Firewall, DNS snooping, and so on). UPF is currently designed to offer Integrated Deep Packet Based Inspection (DPI) Services.

A single instance of UPF provides some or all the following functionalities:

- Anchor point for Intra-RAT and Inter-RAT mobility (when applicable).
- External PDU session point of interconnect to Data Network.
- Packet routing and forwarding.
- Packet inspection. For example, Application detection that is based on the service data flow template and the optional PFDs received from the SMF in addition.
- User Plane part of policy rule enforcement. For example, Gating, Redirection, Traffic steering.
- Lawful intercept (UP collection).
- Traffic usage reporting.
- QoS handling for User Plane. For example, Uplink (UL) and Downlink (DL) rate enforcement, Reflective QoS marking in DL, and so on.
- Uplink Traffic verification (SDF to QoS Flow mapping).
- Transport level packet marking in the Uplink and Downlink.
- Downlink packet buffering and Downlink Data Notification triggering.
- Sending and forwarding of one or more "End Marker" to the source NG-RAN node.

The UPF also provides support for an enterprise mobile virtual network operator (MVNO) model, which enables a mobile network operator (MNO) to perform secondary authentication for the leased MVNO subscribers.

Use Cases and Features

Configuration and Deployment Requirement for UPF

With 5G deployment, interoperability is required between Cisco UPF with non-Cisco SMF, and Cisco SMF with non-Cisco UPF. Also, decoupling of configuration-related messaging between SMF and UPF has the following benefits:

- Alignment with 3GPP standards for configuration bifurcation between User Plane and Control Plane.
- Reduced complexity for configuration management on SMF.
- Simplicity and efficiency for the configuration and change management for User Plane related configuration, as it does not require SMF to manage and distribute the configuration.
- Can be enhanced to achieve interworking between non-Cisco SMF and UPFs.

The Cisco UPF supports 3GPP-specified attributes on the N4 interface. In the current architecture, only UPF associates with the SMF.

The following features are related to this use case:

- UPF Deployment Architecture, on page 13
- UPF Local Configuration, on page 561
- N4 Session Management, Node Level, and Reporting Procedures, on page 385
- Session Recovery, on page 483
- 1:1 Redundancy, on page 19
- UPF Ingress Interfaces, on page 547

Anchor Point for Intra-RAT and Inter-RAT Mobility

The UPF is the anchor point between the mobile infrastructure and the Data Network (DN). That is, the encapsulation and decapsulation of GPRS Tunneling Protocol for the User Plane (GTP-U). Intra-RAT mobility like Xn handover and inter-RAT mobility like 4G to 5G and 5G to 4G handover are supported for this use case.

The GTP-U Support, on page 235 feature is related to this use case.

External PDU Session Point of Interconnect to Data Network

The UPF acts as an external PDU session point of interconnect to Data Network and supports N3, N4, and N6 interfaces. The PDU layer corresponds to the PDU that is transported between the UE and the PDN during a PDU session. The PDU session can be of type IPv4 or IPv6 for transporting IP packets. The GPRS tunneling protocol for the user plane (GTP-U) supports multiplexing of the traffic from different PDU sessions by tunneling user data over the N3 interface (between a 5G access node and the UPF) in the core network. The GTP encapsulates all end-user PDUs and provides encapsulation per-PDU session. This layer also transports the marking associated with the QoS flow. The 5G encapsulation layer supports multiplexing the traffic from different PDU sessions over the N9 interface (an interface between different UPFs). It provides encapsulation per PDU session and carries the marking associated with the QoS flows.

The following features are related to this use case:

- Control Plane-Initiated N4 Association Support, on page 97
- N3 Transfer of PDU Session Information, on page 351
- N4 Session Management, Node Level, and Reporting Procedures, on page 385
- Load and Overload Control Over N4/Sx Interface, on page 567

Packet Inspection

The Cisco UPF performs L3/L4 and L7 inspection for the user traffic that is received. L3/L4 inspection involves IP-address/port matching and Deep Packet Inspection involves matching of L7 header fields.

The Deep Packet Inspection and Inline Services, on page 121 feature is related to this use case.

Per QCI packet drop statistics

Table 5: Revision History

Revision Details	Release
Per-QCI Packet Drop Counters	2026.03.0

The Per QCI Packet Drop Statistics feature provides granular visibility into packet loss within the network. It allows operators to:

- Segment packet drop counts based on the QoS Class Identifier (QCI) value.
- Categorize drops into 19 collapsed drop reasons to simplify root-cause analysis.
- Enable per-APN (pdn-instance) statistics visibility for targeted troubleshooting.

This enables operators to correlate packet drops with specific traffic types (such as VoLTE voice, video, or signaling), facilitating proactive monitoring of latency-sensitive traffic.

User Plane Part of Policy Rule Enforcement

Cisco UPF provides different enforcement mechanisms based on policy received from the SMF. The UPF is the boundary between the Access and IP domains and is the ideal location to implement policy-based enforcement. The pcc-rules provided by the PCF and the pre-defined rules on the SMF are uploaded over the N4 interface and installed on the UPF on a per-DNN basis. This allows for dynamic policy changes that enable differentiated charging and QoS enforcement.

- Dynamic and Static PCC Rules, on page 221
- Voice over New Radio, on page 599

Lawful Intercept

Lawful Interception (LI) enables a LEA to perform electronic surveillance on an individual (a target) as authorized by a judicial or administrative order. To facilitate the lawful intercept process, certain legislation and regulations require service providers and Internet service providers to implement their networks to explicitly support authorized electronic surveillance. Actions taken by the service providers include: provisioning the target identity in the network to enable isolation of target communications (separating it from other users' communications), duplicating the communications for the purpose of sending the copy to the LEA, and delivering the Interception Product to the LEA.

For information about the support of Lawful Intercept by UPF, contact your Cisco Account representative.

Traffic Usage Reporting (Charging)

The usage measurement and reporting function in UPF is controlled by the SMF. The SMF controls these functions by:

- Creating the necessary PDRs to represent the service data flow, application, bearer or session (if they are not existing already).
- Creating the URRs for each Charging Key and combination of Charging Key and Service ID. Also, creating URRs for a combination of Charging Key, Sponsor ID, and Application Service Provider ID.

  Please note that, for static rules, the UPF creates the URR ID. The URR ID is created based on the online/offline and Content ID+Service ID combination that is configured on UPF.
- Associating the URRs to the relevant PDRs defined for the PFCP session, for usage reporting at SDF, Session or Application level.
- For online charging, the SMF provisions Volume and Time quota, if it receives it from the Online Charging Server (OCS).

The Charging Support, on page 63 feature is related to this use case.

QoS Handling for User Plane

The 5G QoS model allow classification and differentiation of specific services, based on subscription-related and invocation-related priority mechanisms. These mechanisms provide abilities such as invoking, modifying, maintaining, and releasing QoS Flows with priority, and delivering QoS Flow packets according to the QoS characteristics under network congestion conditions.

The Dynamic and Static PCC Rules, on page 221 feature is related to this use case.

Downlink Packet Buffering and Data Notification Triggering

A Buffering Action Rule (BAR) provides instructions to control the buffering behavior of the UPF. The BAR controls the buffering behavior for all Forwarding Action Rules (FARs) of the Packet Forwarding Control Protocol (PFCP) session. This control is applicable when the PFCP session is set with an Apply Action parameter, which requests packets to be buffered and associated with the respective BAR.

The Idle Mode Buffering and Paging, on page 247 feature is related to this use case.

Forwarding End Markers to the Source NG-RAN Node

At the time of the handover procedure, the PDU session for the UE – which comprises of UPF node – acts as a PDU session anchor and an intermediate UPF terminating N3 reference point. The SMF sends an N4 Session Modification Request message with the new AN Tunnel Info of NG-RAN to specify the UPF to switch to the N3 paths. In addition, the SMF also specifies the UPF to send the End Marker packets on the old N3 user plane path. After the UPF receives the indication, the End Markers are constructed and sent to each N3 GTP-U tunnel toward the source NG-RAN, after sending the last PDU on the old path.

The N4 Session Management, Node Level, and Reporting Procedures, on page 385 feature is related to this use case.

MVNO Support

The UPF provides support for an enterprise MVNO model. A mobile network operator can perform secondary authentication for the leased MVNO subscribers and also support any additional features related to the AAA server.

The following features are related to this use case:

- APN ACL Support, on page 49

  A configurable mechanism to apply traffic classification and policy enforcement on selective subscriber sessions.

- Dynamic and Static PCC Rules, on page 221

  Increase in maximum number of groups per bandwidth policy.

- Virtual Routing and Forwarding, on page 589

  Support for Overlapping IP Pools and IP Pool chunks.

Smart Licensing

The UPF supports Cisco Smart Licensing, a cloud-based licensing model that simplifies the purchase, deployment, and management of Cisco software assets. For more information, see the Smart Licensing collection page.

Deployment Architecture and Interfaces

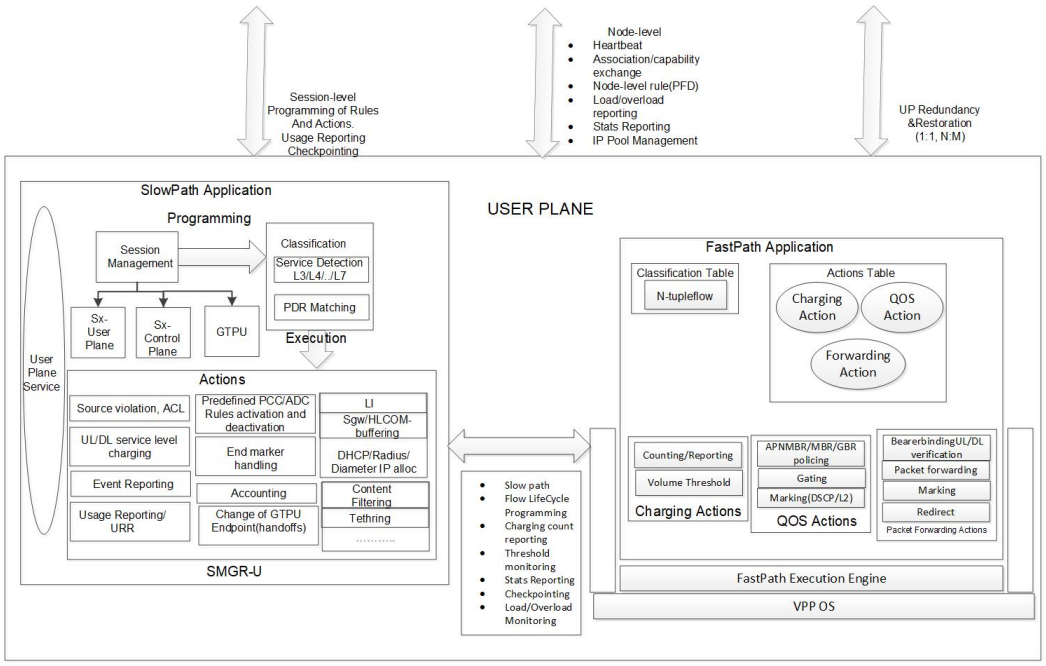
Cisco UPF is part of the 5GC network functions portfolio (AMF/SMF/NRF/PCF/NSSF/UPF) with a common Mobile Core Platform architecture.

UPF Architecture

The User Plane Function (UPF) is a fundamental component of a 3GPP 5G core infrastructure system architecture. The UPF represents the data plane evolution of a Control and User Plane Separation (CUPS) strategy, first introduced as an extension to existing Evolved Packet Cores (EPCs) by the 3GPP in Release 14 specifications. The CUPS decoupless Packet Gateway (P-GW) Control and User Plane functions, enabling the data forwarding component (PGW-U) to be decentralized. This allows packet processing and traffic aggregation to be performed closer to the network edge, increasing bandwidth efficiencies while reducing network load. The P-GW handling signaling traffic (PGW-C) remains in the core, northbound of the Mobility Management Entity (MME).

The primary goal of CUPS is to support 5G New Radio (NR) implementations enabling early IoT applications and higher data rates. Committing to a complete implementation of CUPS is a complex proposition as it only provides a subset of advantages to the operator adopting a 5G User Plane Function (5G-UPF), offering network slicing. Deployed as a Virtual Machine (VM), the User Plane Function delivers the packet processing foundation for Service-Based Architectures (SBAs).

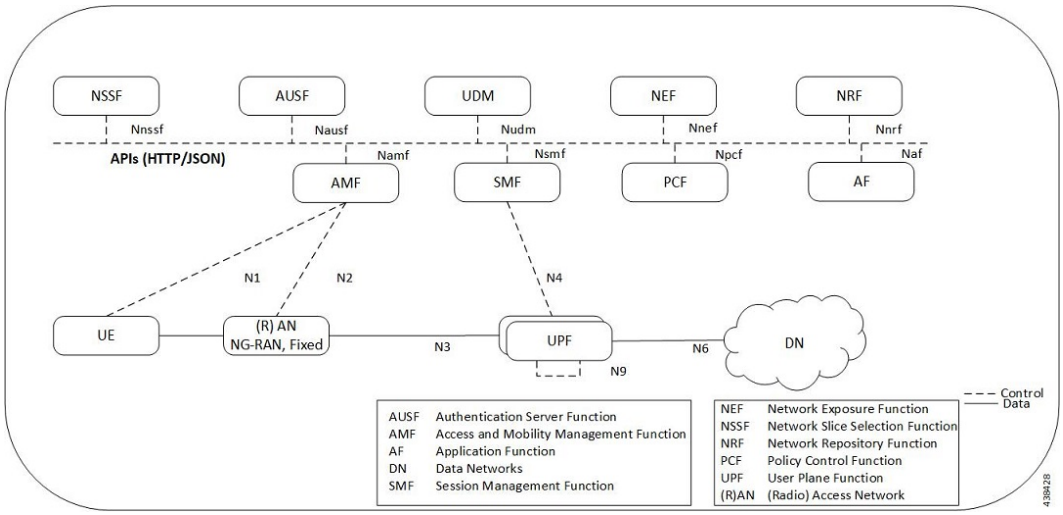
The following figure illustrates the UPF architecture.



The UPF identifies User Plane traffic flow that is based on information that is received from the SMF over the N4 reference point. The N4 interface employs the Packet Forwarding Control Protocol (PFCP), which is defined in the 3GPP technical specification 29.244 for use on Sx/N4 reference points in support of CUPS. The PFCP is similar to OpenFlow but can be limited to only the functionality that is required to support mobile networks. The PFCP sessions, which are established with the UPF, define how packets are identified (Packet Detection Rule / PDR), forwarded (Forwarding Action Rules / FARs), processed (Buffering Action Rules / BARs), marked (QoS Enforcement Rules / QERs) and reported (Usage Reporting Rules / URRs).

UPF Deployment Architecture

The following diagram illustrates the high-level deployment architecture of UPF along with other NFs.



VM Deployment

Virtualized Packet Core—Single Instance (VPC-SI)

VPC-SI consolidates the operations of a physical Cisco ASR 5500 chassis running StarOS into a single Virtual Machine (VM) able to run on commercial off-the-shelf (COTS) servers. VPC-SI can be used as a stand-alone single VM within an enterprise, remote site, or customer data center. Alternatively, VPC-SI can be integrated as part of a larger service provider orchestration solution.

VPC-SI only interacts with supported hypervisors KVM (Kernel-based Virtual Machine) and VMware ESXi. It has little or no knowledge of physical devices.

The UPF functions as user plane node in 5G-based VNF deployments. UPF is deployed as a VNFC running a single, stand-alone instance of the StarOS. Multiple UPF VNFCs can be deployed for scalability based on your deployment requirements.

Hypervisor Requirements

VPC-SI has been qualified to run under the following hypervisors:

- Kernel-based Virtual Machine (KVM) - QEMU emulator 2.0. The VPC-SI StarOS installation build includes a libvirt XML template and ssi_install.sh for VM creation under Ubuntu Server 18.04.

Note When a port on the UPF is shutdown and brought up subsequently, the port interfaces are visible in Ubuntu version 18.04 and NIC driver i40e version 2.12.6. BGP on these interfaces does not recover automatically.

To fully restore the UPF, you must reload the UPFs. In Ubuntu version 20.04 and NIC driver i40e version 2.17.15, both port interfaces and BGP recover automatically.

- KVM - Red Hat Enterprise Linux 7.2: The VPC-SI StarOS installation build includes an install script called qvpc-si_install.sh.
- VMware ESXi 6.7: The VPC-SI StarOS installation build includes OVF (Open Virtualization Format) and OVA (Open Virtual Application) templates for VM creation via the ESXi GUI.

vNIC Options

The supported vNIC options include:

- VMXNET3—Paravirtual NIC for VMware
- VIRTIO—Paravirtual NIC for KMV
- ixgbe—Intel 10-Gigabit NIC virtual function
- enic—Cisco UCS NIC
- SR-IOV—Single-root I/O virtualization

The SR-IOV specification provides a mechanism by which a single root function (for example, a single Ethernet port) can appear to be multiple separate physical devices. Intel 82599 10G is an SR-IOV capable

device and can be configured (usually by the Hypervisor) to appear in the PCI configuration space as multiple functions (PFs and VFs). The virtual functions (VFs) can be assigned to Nova VMs, causing traffic from the VMs to bypass the Hypervisor and go directly to the fabric interconnect. This feature increases traffic throughput to the VM and reduces CPU load on the UCS Servers.

Capacity, CEPS and Throughput

Sizing a VPC-SI instance requires modeling of the expected call model.

Many service types require more resources than others. Packet size, throughput per session, CEPS (Call Events per Second) rate, IPsec usage (site-to-site, subscriber, LI), contention with other VMs, and the underlying hardware type (CPU speed, number of vCPUs) will further limit the effective number of maximum subscribers. Qualification of a call model on equivalent hardware and hypervisor configuration is required.

StarOS VPC-SI Build Components

The following StarOS build filename types are associated with VPC-SI:

- **qvpc-si-<version>.iso** initial installation or startover ISO file.
- **qvpc-si-<version>.bin** update, upgrade, or recovery file for a system that is already running. For additional information, refer to the StarOS Management Operations chapter.
- **qvpc-si-template-libvirt-kvm-<version>.tgz** KVM libvirt template plus ssi_install.sh.
- **qvpc-si-template-vmware-<version>.ova** VMware OVA template.
- **qvpc-si-<version>.qcow2.gz** KVM QCOW2 disk template.

VPC-SI Boot Parameters

The boot parameters file provides a means to pass configuration items to StarOS before it boots. These parameters specify items such as the management, service, and VNFM interface details, and configuration created at StarOS start up.

The boot parameters are sourced in multiple ways, with all methods using the same parameter names and usage. The first location for the boot parameters file is on the first partition of the first VM drive, for example, */boot1/param.cfg*. The second location searched is on the configuration drive, which is a virtual CD-ROM drive. If you are using OpenStack, specify the target boot parameters filename as staros_param.cfg. If you are not using OpenStack, create an ISO image with *staros_param.cfg* in the root directory and attach this ISO to the first virtual CD-ROM drive of the VM.

As the VM boots, the preboot environment known as CFE parses the param.cfg file first. Once the VM starts Linux, the virtual CD-ROM drive is accessed to parse the *staros_param.cfg* file. If there are any conflicts with values that are stored in the */boot1/param.cfg* file, the parameters in *staros_param.cfg* take precedence.

Configuring VPC-SI Boot Parameters

VPC-SI VMs have one interface configured to connect to the management network. This interface is typically configured in StarOS and should be part of the Day 0 configuration. The management interface supports static address assignment through the main StarOS configuration file.

An extra 0 to 4 network interfaces serve as service ports. These interfaces are configured by StarOS. Typically these ports are configured as trunk ports in the VNF infrastructure (VNFI).

VPC-SI VMs have the option of having a network interface that is connected to the virtual network function (VNF) manager (VNFM) if it exists. This interface can be configured through DHCP or static IP assignment and is used to talk to a VNFM or higher-level orchestrator. This interface is enabled before the main application starts.

Sample VPP Configuration

For 5G-UPF, the FORWARDER_TYPE is "vpp".

The following is a sample output of VPP configuration.

```
show cloud configuration
Thursday January 30 12:18:10 UTC 2020
Card 1:
  Config Disk Params:
-------------------------
FORWARDER_TYPE=vpp
VNFM_INTERFACE=MAC:fa:11:3e:22:d8:33
MGMT_INTERFACE=MAC:fa:11:3e:44:af:9e
VNFM_IPV4_ENABLE=true
VNFM_IPV4_DHCP_ENABLE=true
SERVICE1_INTERFACE=MAC:fa:11:3e:11:9d:23
SERVICE2_INTERFACE=MAC:fa:11:3e:99:ec:7b
VPP_CPU_WORKER_CNT=8
VPP_DPDK_TX_QUEUES=9
VPP_DPDK_RX_QUEUES=8
CHASSIS_ID=xxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxx

  Local Params:
-------------------------
    No local param file available
```

Note For additional information about VPC-SI build components, boot parameters, configuring VPC-SI boot parameters, VM configuration, vCPU and vRAM options, VPP configuration parameters, and so on, refer the *VPC-SI System Administration Guide*.

UPF Deployment with VPC-SI

For additional information on VPC-SI, supported operating system and hypervisor packages, platform configurations, software download and installation, and UPF deployment, contact your Cisco Account representative.

For information on Release Package, refer the corresponding Release Notes included with the build.

UPF Deployment with SMI Cluster Manager

The Ultra Cloud Core Subscriber Microservices Infrastructure (SMI) provides a run time environment for deploying and managing Cisco Cloud-Native Network Functions (CNFs), also referred to as applications.

It is built around Open Source projects like Kubernetes (K8s), Docker, Helm, etcd, confd, and gRPC, and provides a common set of services used by deployed cNFs.

The SMI is a layered stack of cloud technologies that enable the rapid deployment of, and seamless life-cycle operations for microservices-based applications.

The SMI stack consists of SMI Cluster Manager that creates the Kubernetes (K8s) cluster and the software repository. The SMI Cluster Manager also provides ongoing Life Cycle Management (LCM) for the cluster including deployment, upgrades, and expansion.

The SMI Cluster Manager leverages the Kernel-based Virtual Machine (KVM)—a virtualization technology—to deploy the User Plane Function (UPF) VMs.

For more information, refer the *UCC SMI Operations Guide*.

Same UP Pools for SAEGW-C and SMF

The same pool of UPs can be used by SAEGW and SMF. The user plane can act as UP and UPF at the same time. It can serve SAEGW over the Sx interface and SMF over the N4 interface. The same subscriber IP pool on SAEGW and SMF is supported only with different VRFs.

This functionality is qualified for the user plane acting as UP and UPF to simultaneously support CUPS and SAEGW Sx interfaces (Sxa, Sxb, and Sxab) for 2G, 3G, 4G RAT, and SMF N4 interface for 5G call.



Note The combined UP and UPF call is not qualified in this release.

Supported Interfaces

This section describes the interfaces supported between the UPF and other network functions in 5GC.

- N3: Interface between the RAN (gNB) and the (initial) UPF; compliant with 3GPP TS 29.281 and 3GPP TS 38.415 (December-2018).
- N4: Interface between the Session Management Function (SMF) and the UPF; compliant with 3GPP TS 29.244 (December-2018).
- N6: Interface between the Data Network (DN) and the UPF; compliant with 3GPP TS 29.561 (December-2018).
- Sx: Interface between the Control-Plane and User-Plane in a split P-GW, S-GW, and TDF architecture in an Evolved Packet Core (EPC); compliant with 3GPP TS 23.214 and 3GPP TS 33.107.

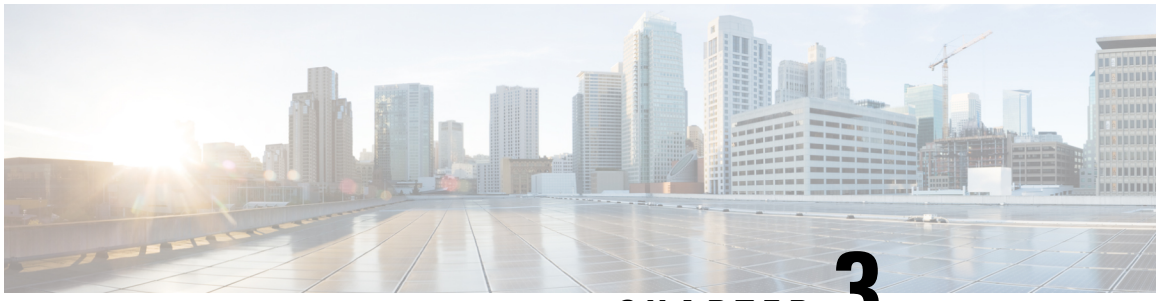
Standards Compliance

Cisco UPF complies with the following standards:

- Interface between the Control Plane and the User Plane Nodes: 3GPP TS 29.244 version 15.4.0. (December-2018)
- General Packet Radio System (GPRS) Tunneling Protocol User Plane (GTPv1-U): 3GPP TS 29.281 version 15.5.0 (December-2018).
- NG-RAN; PDU Session User Plane protocol: 3GPP TS 38.415 (December-2018)
- 5G System; Interworking between 5G Network and external Data Networks; Stage 3: 3GPP TS 29.561 (December-2018)

Limitations

The UPF has the following limitations:



CHAPTER 3

1:1 Redundancy

- Feature Summary and Revision History, on page 19
- Feature Description, on page 20
- How it Works, on page 20
- Configuring 1:1 UPF Redundancy, on page 27
- Monitoring and Troubleshooting, on page 32
- L2 ICSR Support, on page 35

Feature Summary and Revision History

Summary Data

Applicable Product(s) or Functional Area	5G-UPF
Applicable Platform(s)	VPC-SI SMI
Feature Default Setting	Disabled – Configuration Required
Related Changes in this Release	Not Applicable
Related Documentation	Not Applicable

Revision History

Revision Details	Release
Added support for L2 ICSR.	2023.01.0

Revision Details	Release
Support is added for the following functionality: • Zero Accounting Loss in User Plane Function • Early PDU Recovery • Session Prioritization during Recovery • Configuration to change the state of UPF from Pending-Active to Active	2021.02.0
First introduced.	2020.02.0

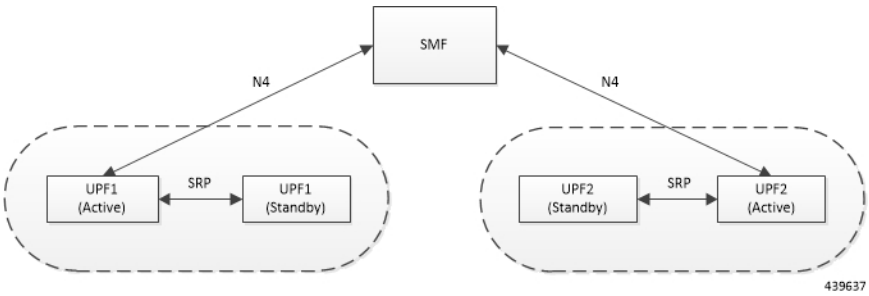
Feature Description

The 1:1 UPF Redundancy feature, for 5G deployment, supports the detection of a failed User Plane Function (UPF) and seamlessly handles the functions of the failed UPF. Each of the Active UPF has a dedicated Standby UPF. The 1:1 UPF Redundancy architecture is based on the UPF to UPF Interchassis Session Recovery (ICSR) connection.

How it Works

The 5G-UPF deployment leverages the ICSR framework infrastructure for checkpointing and switchover of the UPF node as shown in the following figure. The Active UPF communicates to its dedicated Standby UPF through the Service Redundancy Protocol (SRP) link that is provisioned between the UPFs.

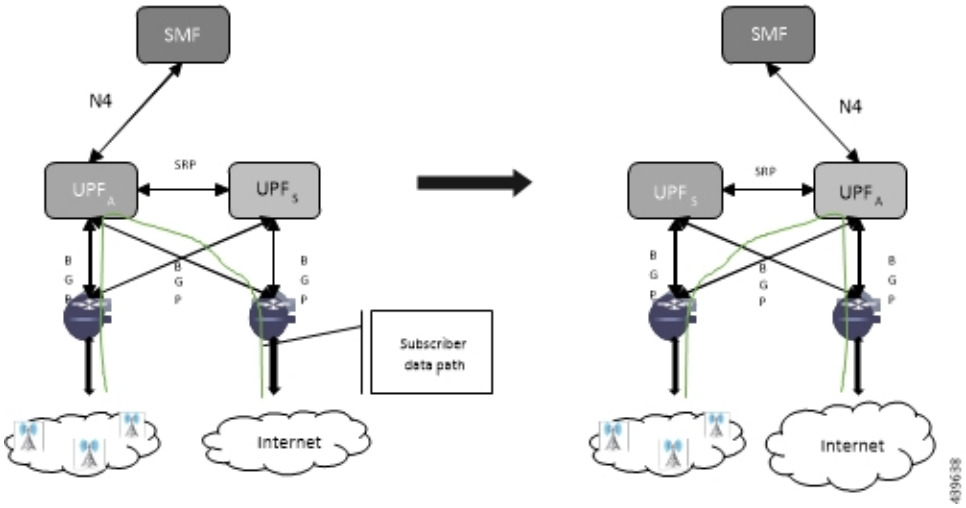
Figure 5: UPF 1:1 Redundancy Using SRP



The Session Management Function (SMF) node does not have the Standby UPF information that is available in the UPF group configuration. Therefore, the SMF is not aware of the UPF redundancy configuration and the switchover event among the UPFs.

The Active UPF communicates to the SMF through the N4 interface address configured in the UPF. The Standby UPF takes over the same Sx/N4 address when it transitions to Active during the switchover event. This implies that the Sx/N4 interface is SRP-activated and is in line with the existing configuration method, therefore UPF switchover is transparent to the SMF.

Figure 6: UPF 1:1 Redundancy Switchover



To make redundancy fully compliant, it addresses the following dependencies on the SRP-based ICSR in the 5G environment.

- Configuration Synchronization (or, Replica Configuration on Standby UPF)
- Sx/N4 Association Checkpoint
- Sx/N4 Link Monitoring

Besides the dependencies listed, the UPF implements data collection and checkpoint procedures specific to the UPF node. For example, checkpointing for IP-pool chunks. The UPF integrates these procedures into the existing ICSR checkpointing framework.

Independent Configuration of Standby UPF

After UPF is up with base configuration (for example, services, contexts, interfaces, and so on), the rest of the configuration (for example, ACS and policy-related configuration) is done through an Ops-center or Redundancy and Configuration Manager (RCM) POD. This configuration is common for both SMF and UPF policies. For SRP redundancy to work, the Active and Standby UPF has same configuration, except SRP-related configuration with which SRP connections are established between the Active and Standby UPF. The RCM configures Active and Standby UPF independently.

BFD Monitor Between Active UP and Standby UP

The Bidirectional Forwarding Detection (BFD) monitors the SRP link between the Active UPF and Standby UPF for a fast failure-detection and switchover. When the Standby UPF detects a BFD failure in this link, it takes over as the Active UPF.

The BFD link can be single-hop or multi-hop.

Note The recommendation is that the SRP bind interface must be an Ethernet interface that attaches to the card service Port. In a loopback address, the recommendation is to ensure that the BFD control packets traverse only through one service port. If it is the ECMP, ensure that the route convergence time does not exceed the BFD timeout.

To configure the BFD monitor, between the Active UP and Standby UP, see *Configuring BFD Monitoring Between Active UPF and Standby UPF*.

Sample Configuration for Multihop BFD Monitoring

Primary UPF:

```
config
  context srp
    bfd-protocol
      bfd multihop-peer 209.165.200.225 interval 50 min_rx 50 multiplier 20
    #exit
    service-redundancy-protocol
      monitor bfd context srp 209.165.200.225 chassis-to-chassis
      peer-ip-address 209.165.200.225
      bind address 209.165.201.1
    #exit
    interface srp
      ip address 209.165.201.1 255.255.255.224
    #exit
    ip route static multihop bfd bfd1 209.165.201.1 209.165.200.225
    ip route 209.165.201.1 255.255.255.224 209.165.201.1 srp
  #exit
end
```

Backup UPF:

```
config
  context srp
    bfd-protocol
      bfd multihop-peer 209.165.201.1 interval 50 min_rx 50 multiplier 20
    #exit
    service-redundancy-protocol
      monitor bfd context srp 209.165.201.1 chassis-to-chassis
      peer-ip-address 209.165.201.1
      bind address 209.165.201.1
    #exit
    interface srp
      ip address 209.165.201.1 255.255.255.224
    #exit
    ip route static multihop bfd bfd1 209.165.200.225 209.165.201.1
    ip route 209.165.201.1 255.255.255.224 209.165.200.225 srp
  #exit
End
```

Router between Primary and Backup UPF:

```
config
  context one
    interface one
      ip address 209.165.201.1 255.255.255.224
    #exit
    interface two
      ip address 209.165.200.225 255.255.255.224
    #exit
```

```
  #exit
end
```

Sample Configuration for Single-Hop BFD Monitoring

Primary UPF:

```
config
  context srp
    bfd-protocol
    #exit
    service-redundancy-protocol
      monitor bfd context srp 209.165.201.1 chassis-to-chassis
      peer-ip-address 209.165.201.1
      bind address 209.165.201.4
    #exit
    interface srp
      ip address 209.165.201.1 255.255.255.224
      bfd interval 50 min_rx 50 multiplier 10
    #exit
    ip route static bfd srp 209.165.201.4
  #exit
end
```

Backup UPF:

```
config
  context srp
    bfd-protocol
    #exit
    service-redundancy-protocol
      monitor bfd context srp 209.165.200.225 chassis-to-chassis
      peer-ip-address 209.165.201.4
      bind address 209.165.201.7
    #exit
    interface srp
      ip address 209.165.201.4 255.255.255.224
      bfd interval 50 min_rx 50 multiplier 10
    #exit
    ip route static bfd srp 209.165.201.7
  #exit
end
```

VPP Monitor

When SRP VPP monitor is configured, the UPF chassis is SRP Active and if the VPP subsystem fails, then SRP initiates switchover to Standby UPF. Currently, VPP health monitoring is limited to heartbeat mechanism between NPUMgr task and VPP process.

Note The VPP monitor is available only on the VPC-SI instance UP. It is not available in the hybrid CUPS ASR 5500 UP because the card level redundancy handles the VPP failure on the ASR 5500. If VPP causes multiple card failures, then SRP card monitor must be used.

To configure the VPP monitor, see *Configuring VPP Monitor on Active UPF and Standby UPF*.

Sx/N4 Association Checkpoint

Whenever an Active UPF initiates an Sx/N4 association to SMF, the Standby UPF checkpoints this data. This maintains the association information even after the UPF switchover.

The Sx/N4 heartbeat messages are sent and the Active UPF responds back even after back-to-back UPF switchovers.

Sx/N4 Monitor

It is critical to monitor the Sx/N4 interface between the UPF and SMF. The SRP monitoring is enabled on Sx/N4 interface and the existing Sx/N4 heartbeat mechanism is leveraged to detect the monitor failure. The Sx/N4 module on Active UPF, on detecting the failure, informs the SRP VPNMgr to trigger UPF switchover event so that the Standby UPF takes over.



Note Sx/N4 monitoring is available only in the UPF.

It is important to ensure that the SMF Sx/N4 heartbeat timeout is higher than the UPF Sx/N4 heartbeat timeout plus UPF ICSR switchover time. This is to ensure that the SMF does not detect the Sx/N4 path failure during a UPF switchover because of the UPF Sx/N4 monitor failure.

The Standby UPF itself has no independent connectivity to the SMF. The Active UPF Sx/N4 context is replicated to the Standby UPF so that it is ready to takeover during SRP switchover. This implies that when the Active UPF has switched over to Standby because of Sx/N4 monitor failure, the new Standby has no way of knowing if the UPF to SMF link is working. To prevent a switchback of the new Standby to Active state again due to Sx/N4 monitor failure in new Active, use the **disallow-switchover-on-peer-monitor-fail** keyword in the **monitor sx** CLI command.

After a chassis becomes Standby due to Sx/N4 monitoring failure, the Sx/N4 failure status is not reset even if Sx/N4 up checkpoint is received from the new Active UPF. This is to prevent the new Active to cause an unplanned switchback again due to Sx/N4 monitor failure when the previous cause of switchover itself was Sx/N4 monitor failure. This prevents back-to-back switchovers when SMF is down. The Sx/N4 monitor failure status must be manually reset when the operator is convinced that the network connectivity is normal. To reset, use the new **srp reset-sx-fail** CLI command (see *Resetting Sx/N4 Monitor Failure*) in the Standby chassis.

To configure the Sx/N4 monitor, see *Configuring Sx/N4 Monitoring on the Active UPF and Standby UPF*.

Sx/N4 Monitor—Pending-Active

The UPF chassis can turn into Pending-Active state for one of the following reasons:

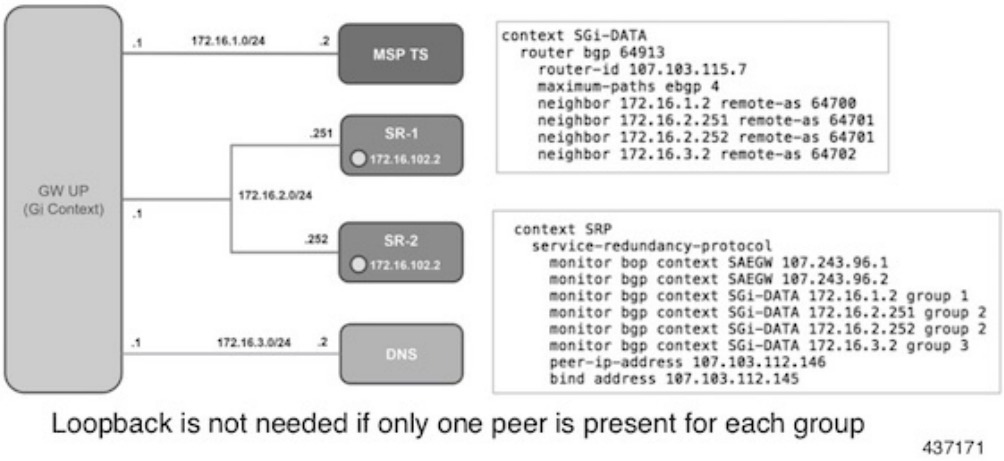
- When Sx/N4 heartbeat times out during SMF upgrade, the Sx/N4 connection is terminated. So, Sx/N4 monitoring failure triggers ICSR switchover in UPF. This switchover causes the old Standby UPF to transition to Pending-Active state. The UPF in Pending-Active state neither receives any Sx/N4 heartbeats from SMF nor any subscriber traffic. As a result, the UPF remains in Pending-Active state indefinitely and can't be utilized without a manual intervention.
- When appropriate procedure to upgrade UPF is not followed, one of the UPF may end up in Pending-Active state. Also, if SMF goes down during the UPF upgrade or if the UPF switchover takes more time than the SMF heartbeat timeout, then one of the UPF remains in Pending-Active state indefinitely.
- When Sx/N4 session times out between SMF and UPF due to network issues, and if a UPF ICSR switchover happens almost simultaneously (Double fault scenario), the UPF in Pending-Active state doesn't transition to active state.

Whenever a UPF chassis turns Pending-Active, start a timer with a callback which forcefully transitions the UPF from Pending-Active to active state. Before forcing the transition, check if the SRP link is up and if the SRP peer is in standby state. If not, restart the timer. The duration of the timer is configurable using **force-pactv-to-actv-timeout** *value_seconds* CLI command (see *Changing UPF state from Pending-Active to Active* section for configuration details). When this CLI command is not configured, the UPF remains in Pending-Active state indefinitely.

BGP Monitor

Configure BGP peer monitor and peer group monitors for the next-hop routers from UPF (both Gi and Gn side). This is the existing ICSR configuration. BGP may run with BFD assist to detect fast BGP peer failure.

Figure 7: BGP Peer Groups and Routing



To configure BGP monitoring and flag BPG monitoring failure, see *Configuring BGP Status Monitoring Between Each UP and Next-Hop Router*.

UPF Session Checkpoints

The Active chassis sends a collection of UPF data as checkpoints to the peer Standby chassis in the following scenarios:

- New call setup
- For every state change in the call
- Periodically for accounting buckets

On receiving these checkpoints, the Standby chassis acts on the data and updates the necessary information either at the call, node, or instance level.

VPN IP Pool Checkpoints

During Sx/N4 Association, the IP pool that is allocated to each of the UPF is sent by SMF to the respective UPF. The VPNMgr receives this message in the UPF and checkpoints the same information to the Standby UPF when the SRP is configured.

The IP pool information is also sent during the SRP VPNMgr restart and during the SRP link down and up scenarios.

Validation of the presence of IP pool information in the Standby is vital before switchover. If the IP pool information is not present, then route advertisement is not possible. Therefore, traffic does not reach the UPF.

External Audit and PFD Configuration Audit Interaction

External Audit management is done in Active UPF. The Session Manager gets a start and complete notification of the Configuration Audit. The Session Manager does not start the External Audit if Configuration Audit is in progress. If the Configuration Audit start-notification arrives when the External Audit is already underway, then the Session Manager raises a flag such that the External Audit restarts when it completes. Restarting the External Audit is necessary because it does not achieve its purpose if it occurs when Configuration Audit is already underway.

Zero Accounting Loss for User Plane Function

Zero accounting loss feature is implemented on the User Plane Function (UPF) so that accounting-data or billing loss is reduced from 18 seconds, which is the default checkpoint time from Active UPF to Standby UPF, or for the configured accounting checkpoint time.

This change in UPF is to support the Gz, Gy, VoGx, and RADIUS URRs. Only planned switchover is supported for zero accounting loss or URR data counters loss. This feature doesn't impact the current ICSR framework or the way checkpointing is done and recovered.

The Sx/N4 usage report is blocked during the “pending active state” until the chassis becomes Active.

Early PDU Recovery for UPF Session Recovery

Early PDU Recovery feature overcomes the earlier limitation of Session Recovery feature wherein it didn't prioritize the CRRs that were selected for recovery. All the CRRs were fetched from the AAAMgr and then the calls were recovered sequentially. The time taken to fetch all the CRRs was a major factor in the perceived delay during session recovery. When a failure occurred, the delay was sometimes long if there were many sessions in a Session Manager. Also, since the calls were recovered in no particular order, the idle sessions were sometimes recovered before active sessions.



Note The Early PDU Recovery feature can recover a maximum of 5-percent sessions.

Session Prioritization during Recovery

Without this functionality, the Session Recovery function didn't prioritize the sessions selected for recovery and loops through all the calls in the call recovery list, and are recovered sequentially when the session recovery is triggered.

As part of Session Prioritization during Recovery, a separate skip list is maintained only for priority calls so that these records can be sent from AAAMgr immediately without going through the loop, thus leading to quicker recovery of the priority calls and reducing the data outage time.

There are two types of sessions at User Plane—Prioritized sessions and normal sessions. Session is considered to be prioritized session based on message priority flag received from SMF and it's recovered first followed by normal calls. These prioritized sessions also take priority in case of early PDU handling. The early PDU of normal calls initiates recovery only when all prioritized sessions are recovered.

In case of critical flush (GR), checkpoints for prioritized sessions are sent first followed by the normal calls. The data of all the calls (both normal and prioritized) are allowed during switchover.

Note The SMF is responsible to set the priority flags for all the calls. The UPF uses the priority call details that are received from the SMF for the Session Prioritization feature.

Configuring 1:1 UPF Redundancy

The following sections provide information about the CLI commands available in support of the feature.

Configuring BFD Monitoring Between Active UPF and Standby UPF

Use the following configuration to configure Bidirectional Forwarding Detection (BFD) monitoring on the Active UPF and Standby UPF. Configure this command in the SRP Configuration Mode.

```
configure
  context context_name
    service-redundancy-protocol
      [ no ] monitor bfd context context_name { ipv4_address | ipv6_address }
{ chassis-to-chassis | chassis-to-router }
      end
```

NOTES:

- **no**: Disables BFD monitoring on the Active and Standby UPF.
- **context** *context_name* : Specifies the context that is used. It refers to the context where the BFD peer is configured (SRP context).

  context_name must be an existing context expressed as an alphanumeric string of 1 through 79 characters.
- *ipv4 _address* | *ipv6_address*: Defines the IP address of the BFD neighbor to be monitored, entered using IPv4 dotted-decimal or IPv6 colon-separated-hexadecimal notation.

  It refers to the IP address of the configured BFD (ICSR) peer.
- **chassis-to-chassis** | **chassis-to-router**:

  chassis-to-chassis: BFD runs between primary and backup chassis on non-SRP links.

  chassis-to-router: BFD runs between chassis and router.

Caution Don't use the **chassis-to-router** keyword for BFD monitoring on the SRP link between the Active UPF and the Standby UPF.

- This command is disabled by default.

Configuring BGP Status Monitoring Between Each UPF and Next-Hop Router

Use the following commands to configure Border Gateway Protocol (BGP) monitoring between each UPF and next-hop router. The command is configured in the SRP Configuration Mode.

```
configure
   context context_name
      service-redundancy-protocol
         [ no ] monitor bgp context bgp-session-context-name [
nexthop-router-ipv4-address | nexthop-router-ipv6-address ] { vrf
bgp-session-vrf-name } { group group-number  }
         end
```

NOTES:

- **no**: Disables BGP status monitoring on the UPF.
- **bgp context** *bgp-session-context-name*: Specifies the context where BGP peer is configured. *bgp-session-context-name* specifies the context string.
- **nexthop-router-ipv4-address | nexthop-router-ipv6-address**: Specifies the configured BGP peer IPv4 or IPv6 address to monitor.
- **vrf** *bgp-session-vrf-name*: Specifies the BGP VPN Routing and Forwarding (VRF) instance. *bgp-session-vrf-name* specifies the VRF name.
- **group** *group-number* : Specifies the BGP peer group where the BGP peer should be included. *group-number* specifies the group number.

  On implementing this keyword, the behavior is as follows:

  - If any BGP peer in that group is up, the BGP peer group is up.

    Omitting group configuration for a BGP monitor includes that monitor in group 0.

    BGP group 0 monitors in a context from an implicit group. Each context forms a separate BGP group 0 implicit monitor group.

    If any BGP peer group is down, BGP monitor is down.

- This command is disabled by default.

Alternate Algorithm to Flag BGP monitoring failure

In this release, an alternate (new) algorithm is introduced to flag BGP monitoring failure.

Use the following commands to flag BGP monitor failure on a single BGP peer (User Plane Function) failure. This command is configured in the SRP Configuration Mode.

```
configure
   context context_name
      service-redundancy-protocol
         [ no ] monitor bgp exclusive-failover
         end
```

NOTES:

- **no**: Disables flagging of BGP monitor failure on a single BGP peer failure.

- On implementing the new **exclusive-failover** keyword, the behavior is as follows:
  - BGP peer group is Up if any BGP peer in that group is Up.
  - Including a BGP peer in group 0 is same as making it non-group (omitting group).
  - BGP monitor is down if any BGP peer group or any non-group BGP peer is down.
- This command is disabled by default.

Configuring Sx/N4 Monitoring on the Active UPF and Standby UPF

Use the following configuration to configure Sx/N4 monitoring on the Active UPF and Standby UPF. This command is configured in the SRP Configuration Mode.

```
configure
   context context_name
      service-redundancy-protocol
        [ no ] monitor sx [ { context context_name | bind-address {ipv4_address
 |ipv6_address } | { peer-address { ipv4_address | ipv6_address } } ]
         end
```

NOTES:

- **no**: Disables Sx/N4 monitoring on the Active and Standby UPF.
- **context** *context_name* : Specifies the context of the Sx/N4 service.

  context_name must be an existing context expressed as an alphanumeric string of 1 through 79 characters.
- **bind-address {** *ipv4 _address* **|** *ipv6_address* **}**: Defines the service IP address of the Sx/N4 service, entered using IPv4 dotted-decimal or IPv6 colon-separated-hexadecimal notation.

Note The IP address family of the **bind-address** and **peer-address** must be same.

- **peer-address {** *ipv4 _address* **|** *ipv6_address* **}**: Defines the IP address of the Sx/N4 peer, entered using IPv4 dotted-decimal or IPv6 colon-separated-hexadecimal notation.
- **disallow-switchover-on-peer-monitor-fail**:

  Prevents the switchback of the UPF to Active state when the working status of the UPF to SMF link is unknown.
- It's possible to implement this CLI command multiple times for monitoring multiple N4/Sx connections.
- The Sx/N4 monitor state goes down when any of the monitored Sx/N4 connections are down.
- This command is disabled by default.
- A maximum of 18 peer nodes can be configured.

Configuring VPP Monitor on Active UPF and Standby UPF

Use the following configuration to configure Vector Packet Processing (VPP) monitor to trigger UPF switchover on the Active UPF if VPP goes down. This command is configured in the SRP Configuration Mode.

```
configure
  context context_name
    service-redundancy-protocol
      monitor system vpp delay-period seconds
      end
```

NOTES:

- If previously configured, use the **no monitor system vpp** CLI command to disable VPP monitoring on the Active and Standby UPF.
- **vpp delay-period** *seconds* : Specifies the delay period in seconds for a switchover, after a VPP failure. *seconds* must be in the range of 0 through 300.

  If the delay period is a value greater than zero (0), then the switchover is initiated after the specified delay period when VPP fails. The last VPP status notification within the delay period is the final trigger for switchover action. The default value is 0 seconds, which initiates an immediate switchover.

  The need for delay is to address the scenario wherein the VPP is temporarily down and the revival is in process. This implies that a switchover may not be necessary.
- This command is disabled by default.

Preventing User Plane Function Switchback

Use the following configuration to prevent the switchback of the new Standby UPF to Active state again due to Sx/N4 monitor failure in the new Active.

```
configure
  context context_name
    service-redundancy-protocol
      monitor sx disallow-switchover-on-peer-monitor-fail timeout seconds

       end
```

Use either of the following CLIs to allow switchback of the new Standby UPF to Active state.

```
no monitor sx disallow-switchover-on-peer-monitor-fail
```

Or

```
monitor sx disallow-switchover-on-peer-monitor-fail timeout 0
```

NOTES:

- **no**: Disables prevention of switchover.
- **disallow-switchover-on-peer-monitor-fail [timeout** *seconds* **]** : Prevents the switchback of the UPF to Active state when the working status of the UPF to SMF link is unknown.

timeout *seconds*: Timeout after which the switchback is allowed even if the Sx/N4 failure status is not reset in the Standby peer. The valid values range from 0 through 2073600 (24 days).



Note Assigning 0 seconds as the timeout allows unplanned switchover.

If **timeout** keyword is not specified, the Active chassis waits indefinitely for the Sx/N4 failure status to be reset in the Standby peer.

- The default configuration is to allow unplanned switchover due to Sx/N4 monitor failure in all conditions.



Note Manual planned switchover is allowed irrespective of whether this CLI is configured or not.

Preventing Dual Active Error Scenarios

Use the following CLI configuration in CP to prevent dual Active error scenarios for UPF 1:1 redundancy.

```
configure
    user-plane-group group_name
       sx-reassociation disabled
       end
```

NOTE:

- **sx-reassociation disabled**: Disables UP Sx reassociation when the association already exists with the CP.

Resetting Sx/N4 Monitor Failure

Use the following configuration only on the Standby chassis to reset the Service Redundancy Protocol (SRP) Sx/N4 monitor failure information. This command is configured in the Exec Mode.

```
srp reset-sx-fail
```

Changing UPF State from Pending-Active to Active

Use the following configuration to change the UPF chassis state from Pending-Active to Active.

```
configure
  context context_name
    service-redundancy-protocol
      force-pactv-to-actv-timeout value_seconds
```

NOTES:

- *value_seconds*: Specifies the timeout value in seconds and must be in the range of 1-300.
- Use the **show config context** *context_name* CLI command to verify the configuration.

Monitoring and Troubleshooting

This section provides information regarding the CLI command available in support of monitoring and troubleshooting the feature.

Show Command(s) and/or Outputs

This section provides information regarding show commands and/or their outputs in support of this feature.

show srp monitor bfd

The output of this CLI command contains the following fields for the 5G UPF 1:1 Redundancy feature:

- Type:
  - (A) - Auth. probe
  - (B) - BGP
  - (D) - Diameter
  - (F) - BFD
  - (E) - EGQC
  - (C) - Card
  - (V) - VPP
- State:
  - (I) - Initializing
  - (U) - Up
  - (D) - Down
- GroupId
- IP Addr
- Port
- Context (VRF Name)
- Last Update

show srp monitor bgp

The output of this CLI command contains the following fields for the 5G UPF 1:1 UPF Redundancy feature:

- Type:
  - (A) - Auth. probe
  - (B) - BGP

- (D) - Diameter
- (F) - BFD
- (E) - EGQC
- (C) - Card
- (V) - VPP
- (S) - Sx

- State:
  - (I) - Initializing
  - (U) - Up
  - (D) - Down

- GroupId
- IP Addr
- Port
- Context (VRF Name)
- Last Update

show srp monitor sx

The output of this CLI command contains the following fields in support of Sx/N4 monitor status:

- Type:
  - (A) - Auth. probe
  - (B) - BGP
  - (D) - Diameter
  - (F) - BFD
  - (E) - EGQC
  - (C) - Card
  - (V) – VPP
  - (S) - SX

- State:
  - (I) - Initializing
  - (U) - Up
  - (D) - Down

- GroupId
- IP Addr
- Port
- Context (VRF Name)
- Last Update

show srp monitor vpp

The output of this CLI command contains the following fields for the 5G UPF 1:1 UPF Redundancy feature:

- Type:
  - (A) - Auth. probe
  - (B) - BGP
  - (D) - Diameter
  - (F) - BFD
  - (E) - EGQC
  - (C) - Card
  - (V) - VPP
- State:
  - (I) - Initializing
  - (U) - Up
  - (D) - Down
- GroupId
- IP Addr
- Port
- Context (VRF Name)
- Last Update

show srp statistics

The output of this CLI command contains the following fields for the Sx/N4 Monitor—Pending-Active functionality:

- Pending-active timer started
- Pending-active timer stopped
- Pending-active to Active forced
- Pending-active to Active force-failed

- Pending-active to Active force-skipped - peer-not-sby
- Pending-active to Active force-skipped - not-PActv

L2 ICSR Support

Feature Description

UPF supports Layer 2-based (L2) ICSR in a two-server architecture to disambiguate packet forwarding between Active and Standby UPFs. UPF uses Gratuitous Address Resolution Protocol (GARP) and Unsolicited Neighbor Advertisement messages to achieve L2 ICSR.

L2 ICSR between SRP peers is configurable in the SRP context. UPF also configures a loopback interface under the same subnet as that of the physical interface. UPF supports SRP BFD monitoring to avoid the Active to Active scenario.

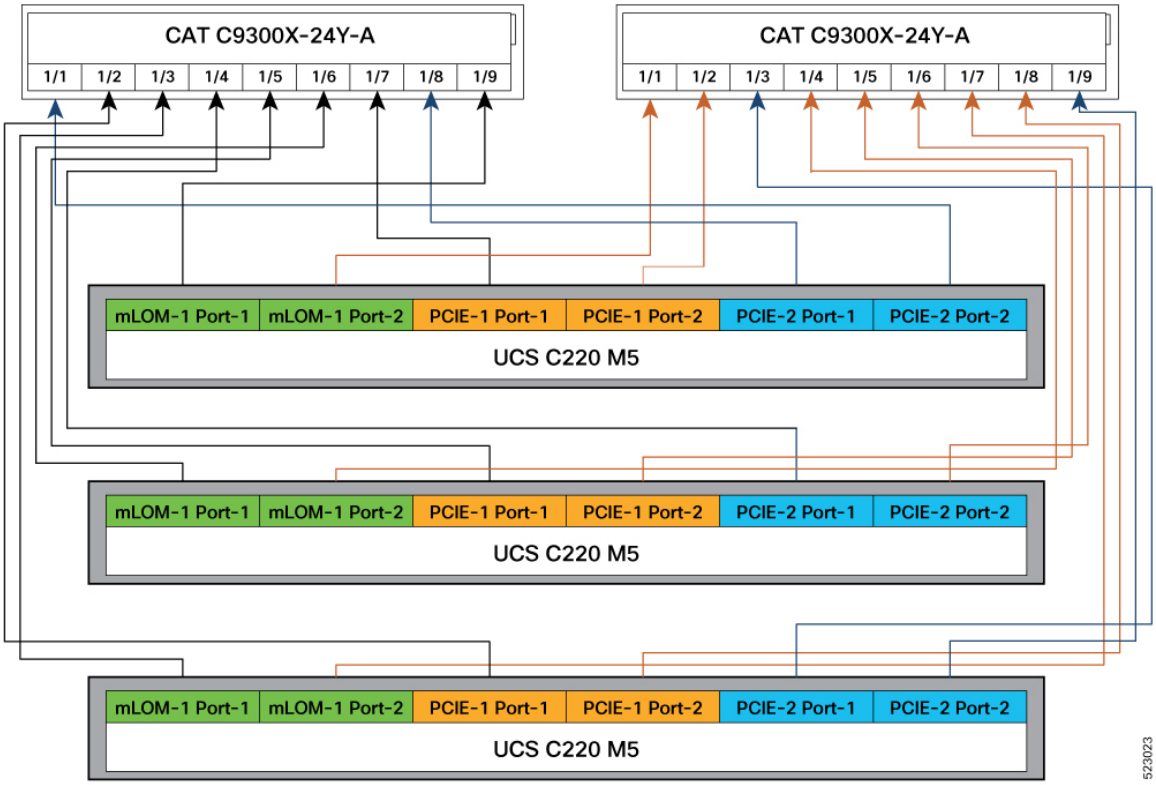
SRP Loopback Interface

SRP loopback interface is the loopback interface for Service IP (N3) or Next-hop IP for IP-Pools (N6). The IP addresses of these loopback interfaces must be on the same subnet as the ethernet interfaces of the context. The SRP loopback interfaces have the same IP address on both Active and Standby UPFs.

- When UPF is Standby, SRP loopback interfaces will not send ARP reply or GARP packets. For IPv6, the UPF does not respond with the Solicited or the Unsolicited Neighbor Advertisement message.
- When UPF is Standby, SRP loopback interfaces are enabled with an IP address attached inside boxer so that the UPF services bound to these interfaces remain in running state.
- When UPF transitions from any other state to the Active state, the SRP loopback interfaces send out GARP packets. For IPv6, the UPF sends Unsolicited Neighbor Advertisement messages.

The following figure illustrates the 3RU wiring diagram for non-bonded interfaces.

Figure 8: 3RU Wiring for Non-Bonded Interfaces



How it Works

UPF directly connects gNB on N3 router or Server on N6 router in the same subnet. UPF ICSR will not be visible to gNB/N3 and Server/N6.

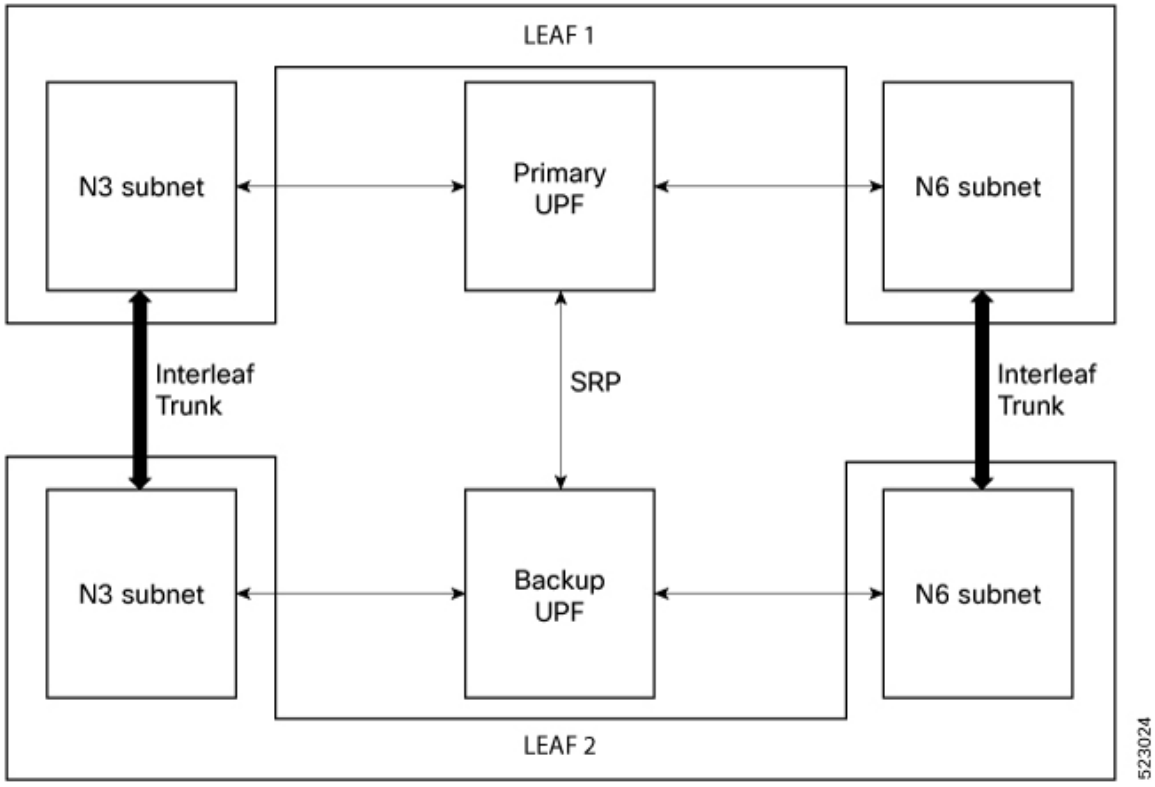
The L2 switch supports BFD. On all N3 and N6 subnets, the switch must have an IP address.

UPF requires Chassis-Chassis (C-R) BFD and Chassis-Router (C-R) BFD for the L2 ICSR solution to work. To avoid the Active-Active state, BFD runs between C-C and C-R. C-C BFD must have a higher timeout value than C-R BFD.

For a two-leaf network design and non-bonded UPF interfaces, each UPF connects to a separate leaf. The IP addresses of the L2 switch are part of the leaf.

The following figure illustrates the high-level diagram of L2 switch wiring for non-bonded interfaces.

Figure 9: L2 Switch Wiring for Non-Bonded Interfaces



State Machine

This table provides information about the UPF states.

Table 6: States Machine

S.No.	Previous State	Internal Fault	C-R BFD	C-C BFD	SRP Link State	Next State	Pull Down BFD
1	X	Yes	X	X	X	Standby	Yes
2	X	X	Any-Down	X	X	Standby	Yes
3	X	No	All-Up	Any-Down	X	Active	No
4	X	No	All-Up	All-Up	Down	No-Change	No
5	X	No	All-Up	All-Up	Up	SRP-Protocol -Decided	No

NOTES:

- For C-R BFDs going from "Not-All-Up" to "All-Up", a delay is introduced in SRP State processing to allow all C-C BFDs to reach steady state.

- To avoid the dual-active scenario, ensure that the states in row 3 of the above table do not occur in both UPFs. This must not happen in a properly functioning and configured system.

Configuring L2 ICSR in SRP Context

To enable or disable L2 ICSR between SRP peers, use the following configuration:

```
configure
  context context_name
    service-redundancy-protocol
      l2-icsr [ disable | enable ]
      end
```

NOTES:

- **disable**: Disable L2 ICSR between SRP peers. This is the default setting.
- **enable**: Enable L2 ICSR between SRP peers.

Configuring C-C and C-R BFD from Primary and Backup Chassis

To configure the C-C and C-R BFD from a Primary and Backup Chassis use the following configuration. The C-C BFD configures as multi-hop to give a different timeout value than the C-R BFD.

```
context SAEGW
bfd-protocol
  bfd multihop-peer 209.165.200.227 interval 50 min_rx 50 multiplier 6
#exit
interface S5_SGW_PGW
  ip address 209.165.200.226 255.255.255.0
  ipv6 address 2001:DB8:A:B::1/64 secondary
  bfd interval 50 min_rx 50 multiplier 3
#exit
ip route static bfd S5_SGW_PGW 209.165.200.225
ip route static multihop bfd bfdcc 209.165.200.226 209.165.200.227
ip igmp profile default
#exit
```

Verifying L2 ICSR Configuration

To verify the configuration, use the following commands:

- **show config context srp**
- **show srp monitor all**
- **show srp info**

The following is a sample output of **show config context srp** command:

```
[local]upf# show config context srp
    context srp
    bfd-protocol
    #exit
    service-redundancy-protocol
      checkpoint session duration non-ims-session 0
      checkpoint session duration ims-session 0
      force-pactv-to-actv-timeout 1
```

```
      l2-icsr enable
      [output truncated]
```

The following is a sample output of **show srp monitor all** command:

```
[local]upf# show srp monitor all
+----- Type: (A) - Auth. probe (B) - BGP (D) - Diameter (F) - BFD (E) - EGQC (C) - Card
| (V) - VPP (S) - Sx
|
|+---- State: (I) - Initializing (U) - Up (D) - Down
||
||+--- GroupId
|||
vvv IP Addr Port Context(VRF Name) Last Update
-------------------------------------------------------------------------------
FU0 209.165.200.225 3785 SAEGW Thu Jan 05 01:13:08 2023
FU0 209.165.200.227 3785 SAEGW Thu Jan 05 01:13:08 2023
FU0 209.165.200.230 3785 ISP Thu Jan 05 01:13:08 2023
FU0 209.165.200.231 3785 ISP Thu Jan 05 01:13:08 2023
-------------------------------------------------------------------------------
[local]upf#
[local]upf# context SAEGW
[SAEGW]upf# show bfd neighbors >>>>>
All the time intervals are in ms
OurAddr NeighAddr LD/RD DtctTime State Intf
209.165.200.226 209.165.200.225 1/1090519046 150 Up S5_SGW_PGW
209.165.200.226 209.165.200.227 2/1 300 Up
Total Sessions 2
Sessions in UP state 2
```

The following is a sample output of **show srp info** command:

```
sample o/p:
[local]l10-ru34-up1# show srp info
Service Redundancy Protocol:
-------------------------------------------------------------------------------
Context: srp
Local Address: 209.165.201.234
Chassis State: Active
Chassis Mode: Primary
Chassis Priority: 47299
Local Tiebreaker: 00-50-56-AE-EE-8D
Route-Modifier: 33

Peer Remote Address: 209.165.201.235
Peer State: Standby
Peer Mode: Backup
Peer Priority: 1673325643
Peer Tiebreaker: 00-00-00-00-00-00
Peer Route-Modifier: 0
Last Hello Message received: Tue Jan 10 04:40:43 2023 (2 seconds ago)
Peer Configuration Validation: Complete
Last Peer Configuration Error: None
Last Peer Configuration Event: Tue Jan 10 04:32:31 2023 (494 seconds ago)
Last Validate Switchover Status: None
Connection State: Connected
Next Peer Audit Scheduled:
Periodicity Audit Scheduled : -
Daily-Start-time Audit Scheduled : -
Cron-Job Audit Scheduled : -
Peer Audit State: Unknown
Last Peer Audit Type: Unknown
Last Peer Audit Successful: -
DSCP Configuration:
Control flow : be
```

```
Session flow : be
Feature Configured-status Operational-status
allow-volte-data-traffic: Disabled Enabled
allow-all-data-traffic: Disabled Enabled

[local]l10-ru34-up1#
```

OAM Support

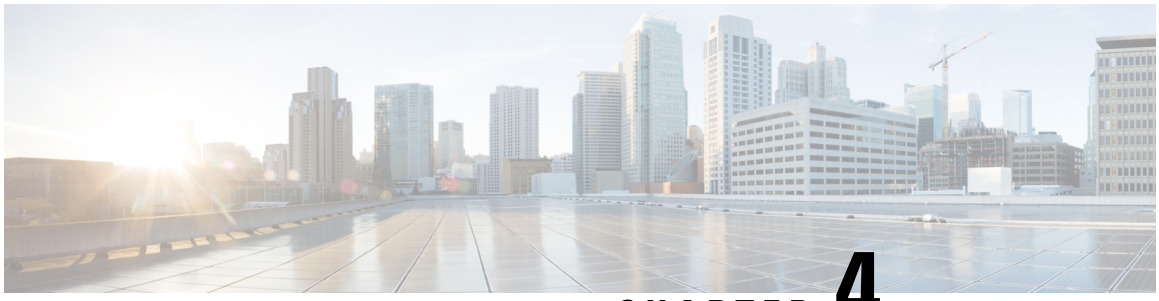
This section describes operations, administration, and maintenance information for this feature.

Show Command Support

Use the **show ip interface** or **show ipv6 interface** command to display the GARP statistics.

The following is a sample output of this command:

```
[EPC2]l10-ru32-up1# show ip interface
Interface Name                   : S5_SGW_PGW
Interface Type                   : Broadcast
Description                      :
VRF                              : None
IP State                         : UP (Bound to 1/10 vlan id 2141, 802.1P prior 0, ifIndex
 17432595)
IP Address                       : 209.165.200.226
Subnet Mask                      : 255.255.255.0
Bcast Address                    : 209.165.200.255
MTU                              : 1500
Resoln Type                      : ARP
ARP timeout                      : 1200 secs
L3 monitor LC-port switchover    : Disabled
Number of Secondary Addresses    : 1
    IPv6 Address: 2001:DB8:A:B::1/64
GARP sent through this interface : 45
GARP sent for this interface     : 0
GARP sent for/through this interface successfully: 45
Last GARP sent for/through this interface: Tue Dec 06 18:40:26 2022
```



CHAPTER 4

ADC for Dynamic Rules

Table 7: Feature History

Feature Name	Release Information	Description
IoT Application Traffic Prioritization with ToS Values	2025.01.0	This feature allows IoT application traffic to be prioritized using ToS values received from cnPGW through dynamic rules. The UPF applies these values to manage data packet traffic efficiently. These changes are visible in the **show subscribers user-plane-only full all** command. **Default Setting**: Enabled – Always-on
ADC Dynamic Rules over Gx	2024.02.0	UPF leverages the dynamic ADC rules for traffic matching and charging. This feature allows the service providers to manage IoT devices, such as connected cars, and charge their subscribers based on the traffic flows classified by SMF or UPF. With this traffic classification, the service providers enable service monetization. In summary, SMF first processes the dynamic ADC rules received from PCF with TDF-App-Identifier, Service ID, and Rating group. SMF then sends this information to UPF to classify ruledefs and perform charging. **Default Setting**: Enabled – Always-on

- Feature Description, on page 41
- How it Works, on page 42
- Limitations, on page 43
- OAM Support, on page 43

Feature Description

UPF processes N4 messages received from SMF and installs, removes, and modifies the message entities (PDR/URR/FAR/QER) of dynamic ADC rules.

UPF matches the configured ruledef with the TDF-App-Identifier AVP received in the dynamic ADC rule as part of PDR installation from SMF. If matched, UPF classifies the packets and charges the subscribers as per the Rating Group and Service ID of the Dynamic rule PDR.

The TosTrafficClass field is received in a dynamic rule from PCRF to SMF and to UPF. It supports the prioritization of traffic flow during network congestion.

How it Works

PCRF installs a dynamic ADC rule with charging rule definition on SMF. UPF receives the N4 messages from SMF containing the dynamic ADC rule-related PDRs, URRs, QERs, and FARs. The dynamic ADC rule PDR contains the Application ID corresponding to the TDF-App-Identifier received from the PCRF and the Mute notification. UPF does not enable any Application Start/Stop triggers as Mute is enabled.

Upon receiving the N4 message, UPF matches the ruledef pointed using the Application ID/TDF-App-Identifier. UPF then charges the packets matching to this ruledef against the URR received and applies bandwidth policing based on the QER received for the dynamic ADC rule PDR.

UPF supports the following three use cases to support dynamic ADC rules:

- **Installing a Dynamic ADC Rule per Subscription**: Upon receiving the dynamic ADC rule from PCRF, SMF installs the rule and sends an N4 Session Establishment Request to UPF to create PDR, QER, FAR, and URR. The Create PDR Request for this dynamic ADC rule contains the TDF-App-Identifier and Mute AVP enabled for all dynamic ADC rules.

  The ToSTrafficClass is an optional value received in SMF through CCA-I message.

  The TOS value and the corresponding mask is populated in TRANSPORT LEVEL MARKING IE in DL FAR and in INNER PACKET MARKING IE in UL FAR respectively in N4 Establishment request.

  UPF matches the configured ruledef with the TDF-App-Identifier and charges the packets according to the RG/SI received in the Dynamic ADC rule. UPF does not generate the N4 Session report for APP-START/STOP, as Mute notifications are enabled from PCRF on SMF during rule install.

- **Modifying a Dynamic ADC Rule from PCRF/SMF**: PCRF sends a Re-authorization request (RAR) to SMF along with a Dynamic ADC Rule Modification request. SMF then sends a Session Modification Request to UPF to update QER containing the updated gate status and MBR values.

  UPF performs modifications of the AVPs and sends a N4 Session Modification Response to SMF. UPF also allows enabling or disabling the offline/online charging and modification of the QoS through RAR in mid-session.

Note Modification of TOS for an existing ADC Rule is not applicable. When ADC Rules are getting installed, they can get installed with different TOS values, however once the installation is successful, the TOS value of the rule will not be changed during the lifetime of a call.

- **Removing a Dynamic ADC Rule from PCRF/SMF**: UPF receives an N4 Modify Request with Remove PDRs and QERs corresponding to the dynamic ADC rule. UPF also receives Remove URRs occasionally, if it is the last rule for that Rating-Group/Service-ID.

  Upon receiving the N4 Modify Request, UPF stops matching the rule for the ruledef pointed by the TDF-App-Identifier from the Session Modification Request.

For detailed information on dynamic ADC rules, see the *UCC 5G SMF Configuration and Administration Guide* applicable for the release.

Limitations

Following is the known limitation of this feature:

- Default bearer URR does not display usage for the subscriber when offline charging is enabled using the RAR on mid-session for the dynamic ADC rule.

OAM Support

This section describes operations, administration, and maintenance support for this feature.

Show Command Output

This section describes the show commands supported in this feature.

show subscribers user-plane-only full all

The output of the show command **show subscribers user-plane-only full all** has been enhanced to display the number of associated ADC PDRs and ToS counters.

```
[local]UPF1# show subscribers user-plane-only full all

  Local SEID      : [0x0004000000000002] 1125899906842626
  Remote SEID     : [0x00000004bebf6665] 20380083813
  State           : Connected
  Connect Time    : Wed Apr  3 12:45:51 2024
  Idle time       : 00h00m07s
  Access Type: uplane-ipv4v6             Network Type: IP
  user-plane-service-name: user-plane-service
  active-service-scheme-name:
  Callid: 00004e23
  Rulebase: adc1
  Interface Type: N4
  eMPS Session: No
  eMPS Session Priority: 0
  Session-Type: normal
  Precedence-order: 0
  Data Queue: 0
  Card/Cpu: 1/0                          Sessmgr Instance: 1
  IP address: 2001:cb0:0:3::4e:2301,11.11.0.4
  Next Hop Ip Address: N/A,
  Source context: EPC2-UP                Destination context: ISP
  PDN-Instance: cnpgw
  User-plane-Sx-addr: 20.20.20.106
  Control-plane-Sx-addr: 20.20.20.3
  Number of associated PDRs : 4
  Number of associated ADC PDRs : 2
  Number of associated FARs : 4
  Number of associated QERs : 3
  Number of associated BARs : 0
  Number of associated URRs : 5
  Number of associated BLIs : 1
  Uplink APN AMBR (bps) : 100000       Downlink APN AMBR (bps) : 100000
  active input acl: n/a                active output acl: n/a
  active input ipv6 acl: n/a             active output ipv6 acl: n/a
```

```
 Bandwidth Policy: n/a
 FW-and-NAT Policy: n/a
 FW-and-NAT Policy ID: n/a
 Firewall Policy IPv4: n/a
 Firewall Policy IPv6: n/a
 NAT Policy NAT44: n/a
 NAT Policy NAT64: n/a
 Converged Session: No                  Converged Peer Callid:      n/a
 Visited Call: No
 Subscriber Parameters:
 IMSI: 123456781300001
 IMEI: 123456781312381
 MSISDN: 22331101101010
 Rat Type: 6
 Old Rat Type: n/a
 MCC MNC: 123456
 SGSN Address: 20.20.20.15
 ULI: TAI+ECGI=21635400152163540000000d00
 CF Policy ID: n/a
 GX CF Policy ID: n/a
 S-NSSAI: 0x01000002
 Roaming Status:  Homer
 Old Roaming Status: n/a
 input pkts: 5                                  output pkts: 5
 input bytes: 5200                              output bytes: 200
 input bytes dropped: 0                         output bytes dropped: 0
 input pkts dropped: 0                          output pkts dropped: 0
 CF Buffered Uplink Packets: 0                  CF Buffered Downlink Packets: 0
 CF Buffered Uplink Bytes: 0                    CF Buffered Downlink Bytes: 0
 Uplink Packets in Buffer: 0                    Downlink Packets in Buffer: 0
 Buff Over-limit Uplink Pkts: 0                 Buff Over-limit Downlink Pkts: 0
 DDN buffered pkts : 0                          DDN buffered bytes : 0
 DDN buffer overflow drop pkts : 0              DDN buffer overflow drop bytes : 0
 pk rate from user(bps): 0                      pk rate to user(bps): 0
 ave rate from user(bps): 0                     ave rate to user(bps): 0
 sust rate from user(bps): 0                    sust rate to user(bps): 0
 pk rate from user(pps): 0                      pk rate to user(pps): 0
 ave rate from user(pps): 0                     ave rate to user(pps): 0
 sust rate from user(pps): 0                    sust rate to user(pps): 0
 ipv4 bad hdr: 0                                ipv4 ttl exceeded: 0
 ipv4 fragments sent: 0                         ipv4 could not fragment: 0
 ipv4 bad length trim: 0
 input pkts dropped (0 mbr): 0                  output pkts dropped (0 mbr): 0
 ipv4 input acl drop: 0                         ipv4 output acl drop: 0
 ipv6 input acl drop: 0                         ipv6 output acl drop: 0
 ip source violations: 0                        ipv4 output no-flow drop: 0
 ipv6 bad hdr: 0                                ipv6 bad length trim: 0
 ipv4 icmp packets dropped: 0
 APN AMBR Uplink Pkts Drop: 0                    APN AMBR Downlink Pkts Drop: 0
 APN AMBR Uplink Bytes Drop: 0                   APN AMBR Downlink Bytes Drop: 0
 APN AMBR Uplink Pkts IP pref lowered: 0         APN AMBR Downlink Pkts IP pref lowered:
0
 APN AMBR Uplink Bytes IP pref lowered: 0        APN AMBR Downlink Bytes IP pref lowered:
0
 ITC Uplink Pkts Drop: 0                         ITC Downlink Pkts Drop: 0
 ITC Uplink Bytes Drop: 0                        ITC Downlink Bytes Drop: 0
 ITC Uplink Pkts IP pref lowered: 0              ITC Downlink Pkts IP pref lowered: 0
 ITC Uplink Bytes IP pref lowered: 0             ITC Downlink Bytes IP pref lowered: 0
 ITC Terminated Flows:: 0                        ITC Redirected Flows: 0
 Flow Action Terminated Flows: 0                 Flow Action Redirected Flows: 0
 ToS marked Uplink Pkts: 0                       ToS marked Downlink Pkts: 0
 CC Dropped Uplink Pkts: 0                       CC Dropped Downlink Pkts: 0
 CC Dropped Uplink bytes: 0                      CC Dropped Downlink Bytes: 0
 Uplink Inflight Pkts: 0                         Downlink Inflight Pkts: 0
```

```
 QFI Mismatch Uplink Pkts Drop: 0                  QFI Mismatch Downlink Pkts Drop: 0
 Total QoS-Group(s) Active: 0
 DNS-to-EDNS Uplink Pkts: 0                      DNS-to-EDNS Uplink Bytes: 0
 EDNS Response Received: 0

 Flow information:
   Current Active Flows:
     TCP: 1
     UDP: 0
   Total Flows:
     TCP: 1
     UDP: 0
     FP:  1
```

show subscribers user-plane-only callid *callid* pdr full all

The output of the show command **show subscribers user-plane-only callid** *callid* **pdr full all**has been modified to display the following new fields:

- TDF App Id
- TDF Notifications
- Total ADC PDRs found

```
[local]UPF1# show subscribers user-plane-only callid 00004e21 pdr full all

Callid: 00004e21
Interface Type: N4
IP address: 2001:cb0:0:3::4e:2301,11.11.0.4
PDR-ID:                    0x0003
 Rule Name:                port-2000
 TDF App Id:         port-2000
 TDF Notifications:  Disabled
 Hits:                      1 Match Bypassed:             4
 Matched Bytes:           200 Matched Packets:            5
 Precedence:               55
 Source Interface:        Core
 Source Interface Type: N/A
 SDF Filter(s):
  No SDF filters present
 Local F-TEID:            0x0
 Local F-TEID IP Addr:
  IPv4:                  N/A
  IPv6:                  N/A
 Outer header removal:
 Application ID:          N/A
 Associated FARID:        0x80000001
  Destination Interface: Access
  Apply Action:           FORWARD
  Outer Header Creation: GTP-U/UDP/IPv4
  Remote TEID:            0x3000000
  Transport Level Marking:         0X88 0XFC
  Transport Level Marking Options:
    Copy Inner:                    No
    Copy Outer:                    No
  Inner Packet Marking:            N/A
  Remote IP Address:      20.20.20.15
 Remote Port:             N/A
 Associated QERIDs:       2
  QER-IDs                 Correlation-ID
   0x00000003                      n/a
   0x00000001                0x0000000E
```

```
 Associated URRIDs:      4
  URR-IDs                Parent URR-ID
   0x00000011            0x80000002
   0x00000027                 n/a
   0x80000002                 n/a
   0x80000009                 n/a
 Activation Time:        N/A
 Deactivation Time:      N/A
 Bearer Level Info-id:        1
    5QI:                     5
    ARP:                     20
    Charging Id:             4

PDR-ID:                    0x0004
 Rule Name:                port-2000
 TDF App Id:         port-2000
 TDF Notifications:  Disabled
 Hits:                      1 Match Bypassed:             4
 Matched Bytes:          5200 Matched Packets:            5
 Precedence:               55
 Source Interface:     Access
 Source Interface Type: S5U(1)
 SDF Filter(s):
  No SDF filters present
 Local F-TEID:           0x8001
 Local F-TEID IP Addr:
  IPv4:                 20.20.20.46
  IPv6:                 N/A
 Outer header removal:   GTP-U/UDP/IPv4
 Application ID:         N/A
 Associated FARID:       0x80000002
  Destination Interface: Core
  Apply Action:          FORWARD
  Outer Header Creation:
  Remote TEID:           0x0
  Transport Level Marking:         N/A
  Transport Level Marking Options:
    Copy Inner:                    No
    Copy Outer:                    No
  Inner Packet Marking:            0X88 0XFC
  Remote IP Address:
 Remote Port:            N/A
 Associated QERIDs:      2
  QER-IDs                Correlation-ID
   0x00000003                     n/a
   0x00000001                 0x0000000E
 Associated URRIDs:      4
  URR-IDs                Parent URR-ID
   0x00000011            0x80000002
   0x00000027                 n/a
   0x80000002                 n/a
   0x80000009                 n/a
 Activation Time:        N/A
 Deactivation Time:      N/A
 Bearer Level Info-id:        1
    5QI:                     5
    ARP:                     20
    Charging Id:             4

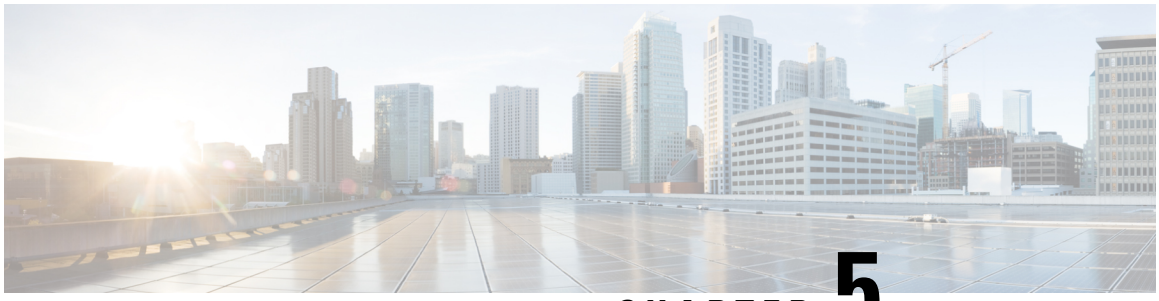
PDR-ID:                    0x0005
 Rule Name:                ipv6rsradl
 TDF App Id:              N/A
 TDF Notifications:       N/A
 Hits:                      3 Match Bypassed:             0
```

```
 Matched Bytes:          312 Matched Packets:            3
 Precedence:               1
 Source Interface:CP-function
 Source Interface Type: N/A
 SDF Filter(s):
  No SDF filters present
 Local F-TEID:           0x8000a001
 Local F-TEID IP Addr:
  IPv4:                 20.20.20.46
  IPv6:                 N/A
 Outer header removal:   GTP-U/UDP/IPv4
 Application ID:         N/A
 Associated FARID:       0x0003
  Destination Interface: Access
  Apply Action:          FORWARD
  Outer Header Creation: GTP-U/UDP/IPv4
  Remote TEID:           0x3000000
  Transport Level Marking:         N/A
  Transport Level Marking Options:
    Copy Inner:                    No
    Copy Outer:                    No
  Inner Packet Marking:            N/A
  Remote IP Address:     20.20.20.15
 Remote Port:            N/A
 Associated QERIDs:      1
  QER-IDs                Correlation-ID
   0x00000002                    n/a
 Associated URRIDs:      0
 Activation Time:        N/A
 Deactivation Time:      N/A
 Bearer Level Info-id:   N/A

PDR-ID:                   0x0006
 Rule Name:               ipv6rsraul
 TDF App Id:             N/A
 TDF Notifications:      N/A
 Hits:                     0 Match Bypassed:             0
 Matched Bytes:            0 Matched Packets:            0
 Precedence:               1
 Source Interface:     Access
 Source Interface Type: S5U(1)
 SDF Filter(s):
  Filter 1
    Protocol:            ICMPV6
    Src IP Addr:         ANY
    Src Port:            ANY
    Dst IP Addr:         ff05::2/128
    Dst Port:            ANY
    SPI:                 0
  Filter 2
    Protocol:            ICMPV6
    Src IP Addr:         ANY
    Src Port:            ANY
    Dst IP Addr:         ff02::2/128
    Dst Port:            ANY
    SPI:                 0
  Filter 3
    Protocol:            ICMPV6
    Src IP Addr:         ANY
    Src Port:            ANY
    Dst IP Addr:         ff01::2/128
    Dst Port:            ANY
    SPI:                 0
 Local F-TEID:           0x8001
```

```
 Local F-TEID IP Addr:
  IPv4:                 20.20.20.46
  IPv6:                 N/A
 Outer header removal:   GTP-U/UDP/IPv4
 Application ID:         N/A
 Associated FARID:       0x0004
  Destination Interface: CP-function
  Apply Action:          FORWARD
  Outer Header Creation: GTP-U/UDP/IPv4
  Remote TEID:           0x4
  Transport Level Marking:         N/A
  Transport Level Marking Options:
    Copy Inner:                    No
    Copy Outer:                    No
  Inner Packet Marking:            N/A
  Remote IP Address:     20.20.20.3
 Remote Port:            N/A
 Associated QERIDs:      0
 Associated URRIDs:      0
 Activation Time:        N/A
 Deactivation Time:      N/A
 Bearer Level Info-id:   N/A

Total ADC PDRs found:  2
Total PDRs found:      4
```



CHAPTER 5

APN ACL Support

This chapter covers the following topics:

- Feature Summary and Revision History, on page 49
- Feature Description, on page 50
- IP Source Violation, on page 53
- Gating Control, on page 53

Feature Summary and Revision History

Summary Data

Table 8: Summary Data

Applicable Product (s) or Functional Area	5G-UPF
Applicable Platforms	VPC-SI SMI
Feature Default Setting	Enabled – Always-on
Related Changes in this Release	Not Applicable
Related Documentation	*UCC 5G UPF Configuration and Administration Guide*

Revision History

Table 9: Revision History

Revision Details	Release
Gating control feature is enhanced to display the gating status as enabled/disabled through an additional QER per PDR for dynamic rules.	2024.01.1

Revision Details	Release
A configurable mechanism has been introduced to apply traffic classification and policy enforcement on selective subscriber sessions.	2021.01.0
First introduced.	2020.02.0

Feature Description

IP Access Lists, commonly known as Access Control Lists (ACLs), control the flow of packets into and out of the system. The configuration is per-context basis and consists of "rules" (ACL rules) or filters that control the action applicable for packets that match the filter criteria. Once configured, an ACL can be applied to an individual subscriber. Separate ACLs can be created for IPv4 and IPv6 access routes.

The following are the two main aspects of ACLs:

- Rule(s)
- Rule Order

Rule(s)

A single ACL consists of one or more ACL rules. Each rule is a filter configured to take a specific action when packets match a specific criteria.

Each rule specifies the action to take when a packet matches the specifies criteria. This section discusses the rule actions and criteria supported by the system.

Actions

ACLs specify that one of the following actions can be taken on a packet that matches the specified criteria:

- **Permit**: The packet is accepted and processed for classification and policy enforcement.
- **Deny**: The packet is rejected.
- **Redirect CSS**: The behaviour is the same as Permit action.

NOTES:

- In UPF, it's recommended to use Permit option instead of Redirect CSS. Functionally, both the options are equivalent in UPF. Support for Redirect CSS option is only for backward compatibility and should be used only in such scenarios.
- Configured ACLs consisting of no rules imply a "deny any" rule. This is the default behavior for an empty ACL.
- In UPF, if ACLs aren't associated with an APN, then call is up. By default, traffic is processed for classification and policy enforcement. For non-UPF architecture, call fails as Redirect CSS is mandatory.
- If only Deny option is given in the ACL for certain traffic, then to pass the rest of the traffic, Permit option must be given explicitly.

- If only permit option is given in the ACL for certain traffic, then to pass the rest of the traffic, permit must be given explicitly for that traffic.
- Router Advertisement/Router Solicitation (RA/RS) packets are candidate for ACL. So, take caution in putting the IPv6 ACL.
- Configuration change in ACL is applied for a new call and not on the existing call.

Criteria

Each ACL consists of one or more rules specifying the criteria that packets will be compared against.

The following criteria are supported:

- **Any**: Filters all packets
- **Host**: Filters packets based on the source host IP address
- **ICMP**: Filters Internet Control Message Protocol (ICMP) packets
- **IP**: Filters Internet Protocol (IP) packets
- **Source IP Address**: Filter packets based on one or more source IP addresses
- **TCP**: Filters Transport Control Protocol (TCP) packets
- **UDP**: Filters User Datagram Protocol (UDP) packets

Each of the above-mentioned criteria is described in detail in the sections that follow.

- **Any**: The rule applies to all packets.
- **Host**: The rule applies to a specific host as determined by its IP address.
- **ICMP**: The rule applies to specific Internet Control Message Protocol (ICMP) packets, Types, or Codes. ICMP type and code definitions can be found at *www.iana.org* (RFC 3232).
- **IP**: The rule applies to specific IP packets or fragments.
- **IP Packet Size Identification Algorithm**: The rule applies to specific IP packets identification for fragmentation during forwarding.

  This configuration is related to the "IP Identification field" assignment algorithm used by the system, when subscriber packets are being encapsulated (such as Mobile IP and other tunneling encapsulation). Within the system, subscriber packet encapsulation is done in a distributed way and a 16-bit IP identification space is divided and distributed to each entity which does the encapsulation, so that unique IP identification value can be assigned for IP headers during encapsulation.

  Since this distributed IP Identification space is small, a nonzero unique identification will be assigned only for those packets which may potentially be fragmented during forwarding (since the IP identification field is only used for reassembly of the fragmented packet). The total size of the IP packet is used to determine the possibility of that packet getting fragmented.
- **Source IP Address**: The rule applies to specific packets originating from a specific source address or a group of source addresses.
- **TCP**: The rule applies to any TCP traffic and could be filtered on any combination of source/destination IP addresses, a specific port number, or a group of port numbers. TCP port numbers definitions can be found at *www.iana.org*.

- **UDP**: The rule applies to any UDP traffic and could be filtered on any combination of source/destination IP addresses, a specific port number, or a group of port numbers. UDP port numbers definitions can be found at *www.iana.org*.

Rule Order

A single ACL can consist of multiple rules. Each packet is compared against each of the ACL rules, in the order in which they were entered, until a match is found. Once a match is identified, all subsequent rules are ignored.

Limitations

Following are the known limitations of APN ACL feature in UPF:

- Readdress option in ACL is not supported.
- Redirect ACL for context and next-hop is not supported.
- Log option is not supported in ACLs.
- APN-level bulkstats for ACL drops (only IPv4) are supported.

Configuring ACL

To apply the ACL to individual subscriber through APN, use the following configuration:

```
configure
   context dest_context_name [ -noconfirm ]
      { ip | ipv6 } access-list acl_list_name
         { permit | deny | redirect }acl
         end
configure
   apn apn_name
      { ip | ipv6 } access-group acl_list_name [ in | out ]
      end
```

Notes:

- The ACL to be applied must be in the destination context of the APN (which can be different from the context where the APN is configured).
- If neither the **in** nor the **out** keyword is specified, the ACL will be applied to all inbound and outbound packets.
- This command supports only one ACL. However, the ACL can have up to 128 rules.
- Four access-groups can be applied for each APN, for example:

  ip access-group acl_list_name_1 in

  ip access-group acl_list_name_2 out

  ipv6 access-group acl_list_name_3 in

  ipv6 access-group acl_list_name_4 out

Verifying ACL Configuration

Use the following CLI commands in Exec mode to check if your ACL lists were applied properly, and also for packet drops due to ACL:

- **show subscriber user-plane-only full all**
- **show subscribers user-plane-only full callid** *call_id*
- **show user-plane-service pdn-instance statistics** *name*

IP Source Violation

Source validation requires the source address of incoming packets to match the IP address of the subscriber during the session. This allows operators to configure the network to prevent problems when a user gets handed back and forth between two gateways several times during a handoff scenario.

When the UPF receives a subscriber packet with a source IP address violation, the system increments the IP source violation drop-limit counter and starts the timer for the IP source violation period. Every subsequent packet that is received with a bad source address during the IP source violation period causes the drop-limit counter to increment. For example, if you set the drop limit to 10, after 10 source violations, the call is dropped. The detection period timer continues to count throughout this process.

The following must be configured in the User Planes APN configuration:

```
ip source-violation { ignore | check [ drop-limit limit ] } [
exclude-from-accounting ]
```

Note For information on IP source violation CLI commands, refer to the StarOS *Command Line Interface Reference*.

Gating Control

Table 10: Feature History

Feature Name	Release Information	Description
Gating Control Using Additional QER	2024.01.1	The number of QERs per PDR supported in UPF is increased from two to three in this release. All rules and filters are sent to UPF through PDRs. SMF sends an additional QER for dynamic rules installed on default QFI/bearer. The Gate Status IE indicates whether the gate is open or closed. **Default Setting:** Not Applicable

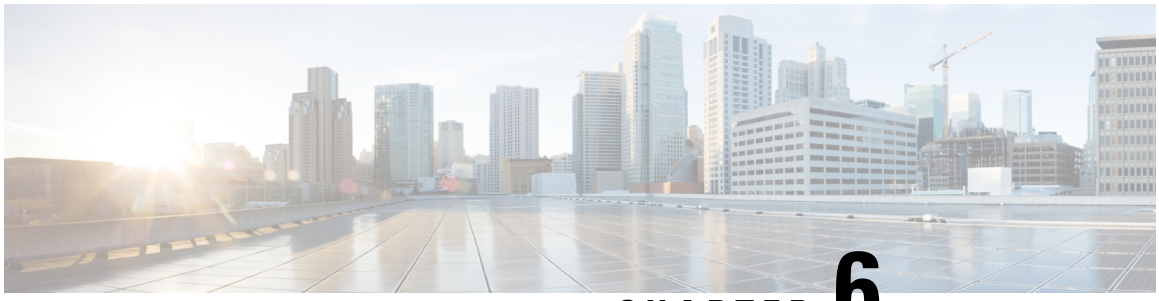
Gating Control in the UPF enables or disables the forwarding of IP packets belonging to a service data flow or detected application's traffic to pass through to the desired endpoint. See 3GPP TS 23.203, subclause 4.3.2.

The SMF controls the gating status in the UPF by creating QERs associated with the PDRs for service data flow(s) or application's traffic to be detected. The QER associated with the Gate Status IE indicates whether the service data flow or detected application traffic is allowed to be forwarded (the gate is open) or to be discarded (the gate is closed) in the uplink and/or in downlink directions.

The default PDRs two QERs get associated to the dynamic rule's PDRs on default QFI/bearer. Therefore, SMF creates an additional QER to display the GATE status for the dynamic rule on default QFI/bearer. The additional QER gets associated with the PDR in the dynamic rule. Hence, the UPF supports up to three QERs per PDR association.

The UPF identifies the UL and DL flows by the Source Interface IE in the PDI of the PDRs or the destination Interface IE in the FARs. The UPF applies UL and DL gating accordingly.

The SMF requests the UPF to discard the packets that are received for the PDR by setting the gate fields in the Gate Status IE of QERs to CLOSED.



CHAPTER 6

APN AMBR Traffic Policing

- Feature Summary and Revision History, on page 55
- Feature Description, on page 56
- Configuring the APN AMBR Traffic Policing Feature, on page 56
- Monitoring and Troubleshooting, on page 57

Feature Summary and Revision History

Summary Data

Table 11: Summary Data

Applicable Product (s) or Functional Area	5G-UPF
Applicable Platforms	VPC-SI SMI
Feature Default Setting	Disabled – Configuration Required
Related Changes in this Release	Not Applicable
Related Documentation	Not Applicable

Revision History

Table 12: Revision History

Revision Details	Release
First introduced.	2021.01.0

Feature Description

The APN-AMBR is a subscription parameter that is stored per APN in the HSS. S-GW provides APN-AMBR during default bearer establishment procedure. APN-AMBR limits the aggregate bit rate that can be expected to be provided across all non-GBR bearers and across all PDN connections of the same APN. Each of those non-GBR bearers can potentially utilize the entire APN-AMBR, for example, when the other non-GBR bearers do not carry any traffic. The P-GW enforces the APN-AMBR in downlink and uplink direction.

Limitations

The **token-replenishment-interval** and **violate-action shape** CLI commands are not supported.

Configuring the APN AMBR Traffic Policing Feature

This section describes how to configure the APN-AMBR Traffic Policing feature.

```
configure
         context context_name
           apn apn_name
       apn-ambr rate-limit direction { downlink | uplink } [ burst-size
{ auto-readjust duration { milliseconds msecs | seconds } | violate-action
 { drop | lower-ip-precedence | transmit }
                end
```

NOTES:

- **rate-limit direction { downlink | uplink }**: Specifies that the rate limit is to be applied to either the downlink (network to subscriber) traffic or the uplink (subscriber to network) traffic.
- **burst-size { auto-readjust duration milliseconds msecs | seconds }**: This parameter is used by policing algorithms to permit short bursts of traffic not to exceed the allowed data rates. It's the maximum size of the token bucket.
  - **auto-readjust duration***seconds*: The duration (in seconds) used in this burst size calculation: burst size = peak data rate/8 * auto-readjust duration.
    - Seconds must be an integer value from 1-30. Default is 1 second.
  - **milliseconds**: *msecs* must be an integer value from 100-900, in increments of 100 milliseconds. For example, 100, 200, or 300, and so on.
- **violate-action { drop | lower-ip-precedence | transmit }**: The action that the P-GW takes when the data rate of the bearer context exceeds the AMBR.
  - **drop**: Drops violating packets.
  - **lower-ip-precedence**: Sets the DSCP value to zero ("best effort") for violating packets.
  - **transmit**: Transmits violating packets. This is the default behavior of the feature.
- Prior to this feature, the default behavior was to drop the violating packets.

Monitoring and Troubleshooting

This section provides information about the commands available to monitor and/or troubleshoot the APN-AMBR Traffic Policing feature.

Show Commands and/or Outputs

This section provides information about the show commands available for monitoring and/or troubleshooting the APN-AMBR Traffic Policing feature.

- **show user-plane-service pdn-instance name <apn_name>**

  Use this show command in UPF to see if the rate limit is enabled/disabled, burst size, and other such parameters for downlink/uplink traffic:

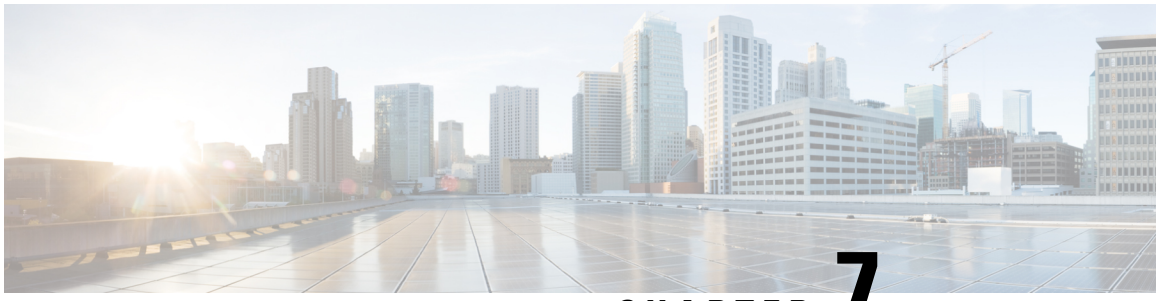
  - APN-AMBR
    - Downlink Apn Ambr: Indicates if the rate limit is enabled or disabled for downlink traffic.
      - Burst Size: Indicates the burst size of the downlink traffic.
      - Auto Readjust: Indicates if the auto-readjust is enabled or disabled for downlink burst size.
      - Auto Readjust Duration: Indicates the duration used in downlink burst size calculation.
      - Burst Size(bytes): Indicates the burst size in bytes.
      - Violate Action: Indicates the action that the P-GW takes when the data rate of the bearer context exceeds the AMBR for downlink traffic.
    - Uplink Apn Ambr: Indicates if the rate limit is enabled or disabled for uplink traffic.
      - Burst Size: Indicates the burst size of the uplink traffic.
      - Auto Readjust: Indicates if the auto-readjust is enabled or disabled for uplink burst size.
      - Auto Readjust Duration: Indicates the duration used in uplink burst size calculation.
      - Burst Size(bytes): Indicates the burst size in bytes.
      - Violate Action: Indicates the action that the P-GW takes when the data rate of the bearer context exceeds the AMBR for uplink traffic.
    - Token Replenishment Interval: Indicates the token replenishment interval duration.

- **show sub user-plane-only full all**

  Use this show command in UPF to see the count of packets that are dropped, and IP precedence lowered due to APN-AMBR policer. The following fields are introduced in support of this feature:

  - APN AMBR Uplink Pkts Drop: Indicates the number of APN-AMBR packets that are dropped for uplink traffic.

- APN AMBR Uplink Bytes Drop: Indicates the number of APN-AMBR bytes that are dropped for uplink traffic.
- APN AMBR Uplink Pkts IP pref lowered: Indicates the number of APN-AMBR uplink packets for which IP precedence is lowered.
- APN AMBR Uplink Bytes IP pref lowered: Indicates the number of APN-AMBR uplink bytes for which IP precedence is lowered.
- APN AMBR Downlink Pkts Drop: Indicates the number of APN-AMBR packets that are dropped for downlink traffic.
- APN AMBR Downlink Bytes Drop: Indicates the number of APN-AMBR bytes that are dropped for downlink traffic.
- APN AMBR Downlink Pkts IP pref lowered: Indicates the number of APN-AMBR downlink packets for which IP precedence is lowered.
- APN AMBR Downlink Bytes IP pref lowered: Indicates the number of APN-AMBR downlink bytes for which IP precedence is lowered.



CHAPTER 7

Bulk Statistics Support

- Feature Summary and Revision History, on page 59
- Feature Description, on page 60
- Bulk statistics for per-QCI drop counters, on page 62

Feature Summary and Revision History

Summary Data

Table 13: Summary Data

Applicable Product (s) or Functional Area	5G-UPF
Applicable Platforms	VPC-SI SMI
Feature Default Setting	Enabled – Always-on
Related Changes in this Release	Not Applicable
Related Documentation	*UCC 5G UPF Configuration and Administration Guide*

Revision History

Table 14: Revision History

Revision Details	Release
The following new bulk statistic schemas are now supported: • P2P Schema • Sx Schema • System Schema • Userplane Schema	2021.01.0
First introduced.	2020.02.0

Feature Description

This chapter identifies bulk statistic schemas for the Cisco Ultra Cloud 5G User Plane Function (UPF) software release.

Bulk statistics is a collection of software features and framework that collects and exports the important performance and health-related statistics of the packet core node to an external node. These statistics provide an effective way for the operators to perform the following functions:

- Monitor the overall health and performance of the nodes.
- Help take corrective actions.
- Optimize the packet core network for better utilization.
- Reduce the overall operation expenses.

The individual statistics are configured to be collected in a group called ‘schema.'

The system-supported bulk statistics allows operators to choose statistics that are of importance to them and configure the presentation format. This simplifies the post-processing of statistical data because it allows data formatting that facilitates external, backend processors to parse it.

Statistics or bulk statistics reporting is important on a Mobile Packet Core node. For a product to be deployed in the network, it has to support statistics that meets Carrier Grade requisites.

Operators use bulk statistics for the following:

- Performance KPI monitoring
- Network Fault analysis and debugging
- Network Optimization
- Traffic pattern analysis
- Node health analysis

When used along with an element management system (EMS), the data can be parsed, archived, and graphed.

In the 5G environment, the system can be configured to collect for the following network functions:

- Access and Mobility Management Function (AMF)
- Network Repository Functions (NRF)
- Network Slice Selection Functions (NSSF)
- Policy Control Function (PCF)
- Session Management Function (SMF)
- User Plane Function (UPF)

The system supports the configuration of up to four sets (primary and secondary) of receivers. Each set is configured to collect specific sets of statistics from the supported list of schemas. Statistics can be pulled manually from the system or sent at configured intervals. The bulk statistics are stored on the receivers in files.

You can configure the format of the bulk statistic data files. You can specify the following:

- Format of the filename
- File headers and footers to include information such as the date, system hostname, and system uptime.
- IP address of the system generating the statistics (available for only for headers and footers)
- Time that the file was generated.

An EMS is capable of further processing the statistics data through XML parsing, archiving, and graphing. The Bulk Statistics Server component of an EMS parses collected statistics and stores the information in its PostgreSQL database. It can also generate XML output and can send it to a Northbound NMS or an alternate bulk statistics server for further processing.

Also, the Bulk Statistics server can archive files to an alternate directory on the server. The directory can be on a local file system or on an NFS-mounted file system on an EMS server.

The implementation of bulk statistics in 5G is as follows:

- The NFs collect and export the statistics separately to an aggregator node in the 5G architecture.
- The receiver correlates the statistics from the NFs using the node-names or any other information that is configured as part of the bulk statistics configuration. Any EMS tool can render this data similar to how it is rendered from a standalone system.

Supported Schemas

This release supports the following schemas in the 5G architecture.

APN Schema

The APN schema provides Access Point Name (APN) statistics.

Card Schema

The Card schema provides card-level statistics.

ECS Schema

The ECS schema provides Enhanced Charging Service statistics.

GTP-U Schema

The GTP-U schema provides GPRS Tunneling Protocol- User message statistics.

P2P Schema

The P2P schema provides point-to-point statistics.

P-GW Schema

The P-GW schema provides user-plane service statistics.

Port Schema

The Port schema provides port-level statistics.

Rulebase

The Rulebase schema provides rule base statistics.

Sx Schema

The Sx schema provides N4 related message statistics.

System Schema

The System schema provides system-level statistics.

Userplane Schema

The Userplane schema provides User Plane statistics.

VLAN-NPU

The VLAN-NPU schema provides VLAN and NPU statistics.



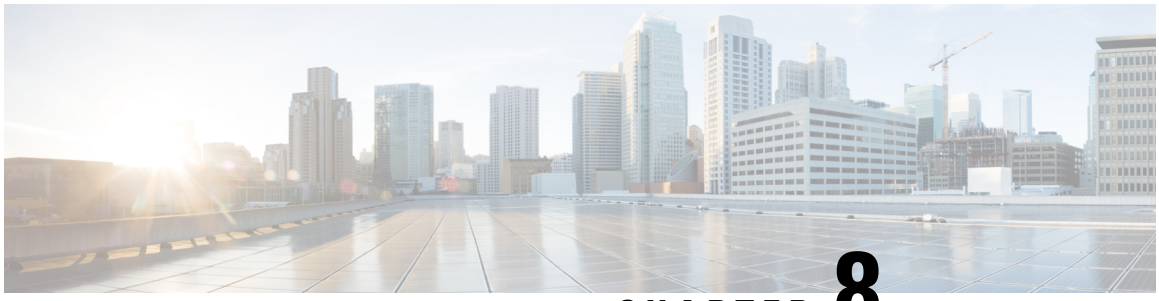
Important For more information on bulk statistic configuration, refer to the *Bulk Statistics* chapter in the *ASR 5500 System Administration Guide*.

Bulk statistics for per-QCI drop counters

A bulk statistics schema is a data collection framework that provides performance and health-related metrics for the 5G User Plane Function (UPF) by:

- Aggregating granular drop counters per QCI and per APN.
- Exporting performance data to external network management systems, and
- Enabling long-term traffic pattern analysis and network optimization.

The system-supported bulk statistics allow operators to choose statistics of importance and configure the presentation format. By collecting per-QCI drop data, operators can identify performance bottlenecks in latency-sensitive traffic and take corrective actions to improve subscriber quality of experience.



CHAPTER 8

Charging Support

This chapter covers the following topics:

- Feature Summary and Revision History, on page 63
- Feature Description, on page 64
- How it Works, on page 65
- Configuring Credit Control for Usage Reporting, on page 75
- Configuring ACS Rulebase for Usage Reporting, on page 76
- Monitoring and Troubleshooting, on page 79

Feature Summary and Revision History

Summary Data

Table 15: Summary Data

Applicable Product(s) or Functional Area	5G-UPF
Applicable Platform(s)	VPC-SI SMI
Feature Default Setting	Disabled - Configuration Required
Related Changes in this Release	Not Applicable
Related Documentation	Not Applicable

Revision History

Table 16: Revision History

Revision Details	Release
The feature is enhanced to support Recalculate IE on N4 interface	2023.04.0

Revision Details	Release
The feature is enhanced to support following functionality: • PTT no-quota Limited Pass • PTT quota exhaust Limited Pass • Tariff Time • TCP Maximum Segment Size	2021.02.0
Usage reporting with Rating-Group and Service ID is introduced.	2020.02.5
First introduced.	2020.02.0

Feature Description

The usage measurement and reporting function in UPF 2.0is controlled by the Session Management Function (SMF). The SMF controls these functions by:

- Creating the necessary PDRs to represent the service data flow, application, bearer or session (if they are not existing already).
- Creating the URRs for each Charging Key and combination of Charging Key and Service ID. Also, creating URRs for a combination of Charging Key, Sponsor ID, and Application Service Provider Id.
- Associating the URRs to the relevant PDRs defined for the PFCP session, for usage reporting at SDF, Session or Application level.
- For online charging, the SMF provisions Volume and Time quota, if it receives it from the Online Charging Server (OCS).

Offline Charging Events Reporting over N4

The User Plane Function (UPF) supports session-based offline charging, PDU session level reporting triggers in URR (volume and time threshold), PFCP session report procedure, and usage report IE support in the PFCP modification response for the Session-AMBR change, QoS, and User Location triggers.

Online Charging Support over N4

The UPF supports flow-based online charging support, which includes URR enhancements for Volume and Time quota and Usage reporting IE in PFCP modify response. In addition, the UPF supports online charging triggers, which include a PFCP session report request support with usage reporting IE.

How it Works

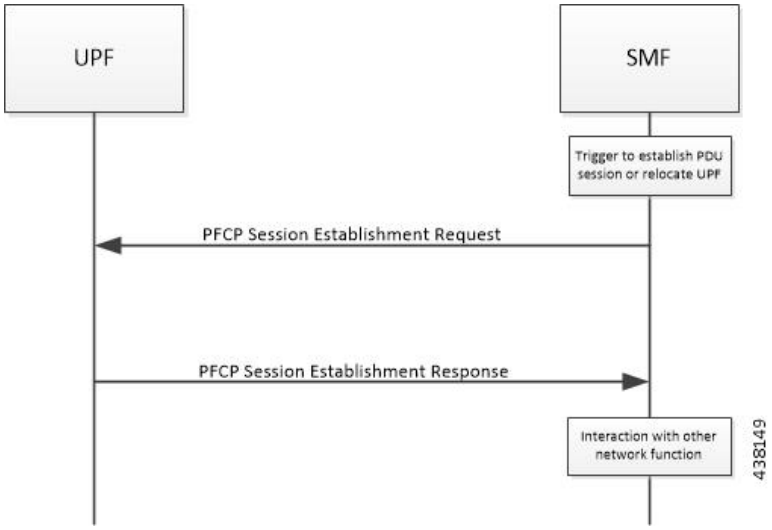
Call Flows

The following sections describe the call flows between SMF and UPF for PFCP Session Management.

PFCP Session Establishment Procedure

The PFCP Session Establishment procedure establishes a PFCP session between SMF and UPF. It also configures rules in UPF for handling incoming packets. In addition, the SMF sends Create URR Information Element (IE), which comprises of triggers and thresholds that are intended for reporting.

The following call flow depicts the PFCP Session Establishment procedure.

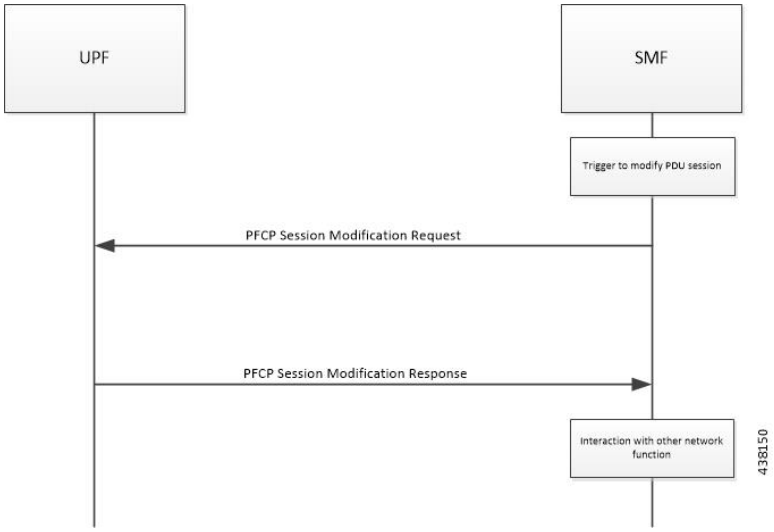


Step	Description
1	The SMF receives the trigger to establish a new PDU Session or change the UPF for an established PDU Session.
2	The SMF sends the PFCP Session Establishment Request message to the UPF. This message contains the structured control information, which defines the UPF's behavior.
3	The SMF provisions URR with Create URR IE. The Create URR associates with PDRs by adding URR-ID IE in Create PDR IE. It includes various triggers and thresholds for usage reporting.
4	When the same URR is associated with multiple PDRs, URRs are linked with another URR. Therefore, if a report for an URR is sent, its linked URR is also reported.
5	The UPF responds with the PFCP Session Establishment Response message to the SMF. For instance, Created PDR IE, in which UPF Flow-TEID is sent to gNB for GTP-u encap for data traffic.
6	The SMF interacts with the network function, which triggered this procedure. For instance, AMF or PCF.

PFCP Session Modification Procedure

The SMF uses the PFCP Session Modification procedure to modify an existing PFCP session on the UPF. For instance, configuring a new rule, modifying an existing rule, or deleting an existing rule, and so on. The SMF sends the Create URR IE, Update URR IE (to update the trigger or threshold) and Remove URR IE (to remove an existing URR created earlier by SMF during Session Establishment Procedure) in the same message.

The following call flow depicts the PFCP Session Modification procedure.

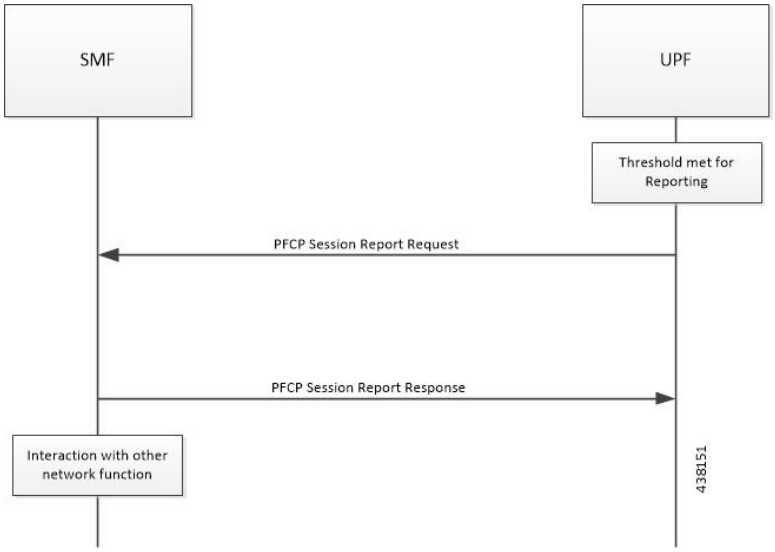


Step	Description
1	The SMF receives the trigger to modify the existing PDU Session.
2	The SMF sends an N4 session modification request message to the UPF. This message contains the structured control information, which defines the UPF's behavior.
3	The UPF identifies the PFCP session context for the Session ID to modify. It updates the parameters of this session context according to the list of parameters sent by the SMF. It then responds with a PFCP Session Modification Response message. The message contains the information, which the UPF must provide to the SMF (in response to the control information received).
4	• If the SMF sends the QAURR flag set in PFCPSMReq-Flag IE or URR ID (s) with Query URR IE (e), then UPF sends the usage report IE for the corresponding URR with the PFCP Session Modification response. • The custom Information Element (IE) called Recalculate Measurement is added to the Update-URR IE (for URR-ID: Gz-Bearer) to support the PGWCDR generation due to the Max-LOSDV condition with the N4 Session Modify Request. After receiving the IE, UPF reconciles the Volume and Duration of the URR according to the URR-ID present inside the IE.
5	The UPF provisions and acts based on the Create URR, Update URR or Remove URR IE sent by the SMF.
6	The SMF interacts with the network function, which triggered this procedure. For instance, AMF or PCF.

PFCP Session Reporting Procedure

The UPF uses PFCP Session Reporting procedure to report information that is related to the PFCP session to the SMF (usage report IE). Once the threshold hits the volume, time or event measurement and sets the corresponding trigger for reporting, the message is sent to the SMF by the UPF.

The following call flow depicts the PFCP Session Reporting procedure.

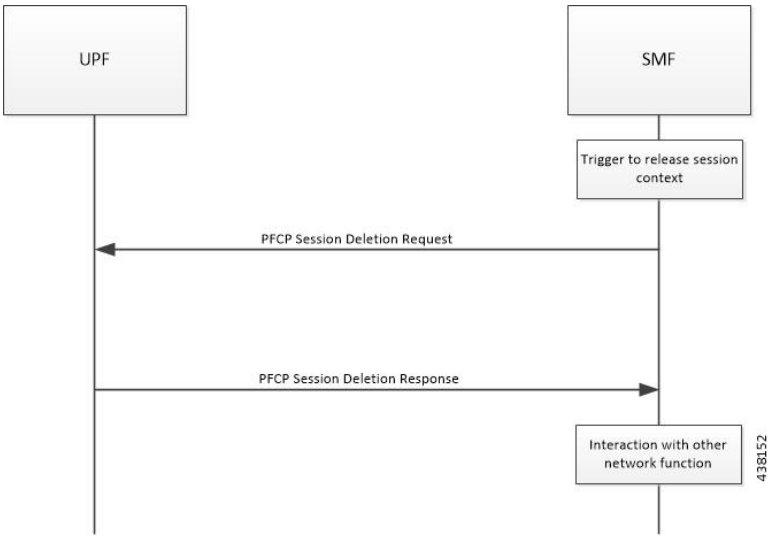


Step	Description
1	• Once the provisioned threshold is met (for time, volume or event, and trigger is set for reporting), the UPF sends PFCP Session Report Request with usage report IE and usage details for volume, time, or threshold. • The custom Information Element (IE) called Recalculate Measurement is added to the Update-URR IE (for URR-ID: Gz-Bearer) to support the PGWCDR generation due to the Max-LOSDV condition with the N4 Session Modify Request. After receiving the IE, UPF reconciles the Volume and Duration of the URR according to the URR-ID present inside the IE.
2	The SMF responds with PFCP Session Modification Response with success or failure message. No failure handling is needed on the UPF.
3	The SMF interacts with the network function, which triggered this procedure. For instance, AMF or PCF.

PFCP Session Deletion Procedure

The PFCP Session Deletion procedure deletes an existing PFCP session between the SMF and UPF. The SMF initiates a PFCP Session Deletion procedure toward the UPF to delete an existing PFCP session. The UPF sends the Session Deletion Response including the Usage Report for all URRs provisioned earlier.

The following call flow depicts the PFCP Session Deletion procedure.



Step	Description
1	The SMF receives the trigger to remove the PFCP session context for the PDU Session.
2	It sends the PFCP Session Delete Request message to the UPF.
3	The UPF identifies the PFCP session context for the Session ID to remove. It then removes the whole session context. In addition, the UPF responds with a PFCP Session Delete Response message that contains any information the UPF provides to the SMF. For instance, the UPF sends usage report for all the URR provisioned for this session.
4	The SMF interacts with the network function, which triggered this procedure. For instance, AMF or PCF.

IEs Supported for Offline Charging Reporting

Table 17: Feature History

Feature Name	Release Information	Description
Recalculate Measurement IE on the N4 Interface	2023.04	UPF supports the Recalculate Measurement custom IE as received over the N4 interface. This IE is added to the Update-URR process (URR-ID: Gz-Bearer) to support the PGW-CDR generation due to the max_LOSDV change condition.

The following trigger Information Elements (IEs) support offline charging Reporting over N4:

- Periodic Reporting – When this trigger is set, the UPF sends resource usage report periodically to Session Management Function (SMF). The intervals that are required for periodic reporting are sent with the measurement period IE.
- Volume Threshold (when the volume threshold reaches UL, DL, and Total) – This trigger is set when the volume-based measurement is required. The SMF sends the traffic volume value along with the

volume threshold IE, while the UPF sends the traffic usage report when the traffic volume is reached for the specific Usage Reporting Rule (URR).

- Time Threshold (when the time threshold is reached) – This trigger is set when the time-based measurement is set. The SMF sends the time threshold value along with the time threshold IE, while the UPF sends resource usage report when the time threshold is reached for the specific URR.
- Linked Usage Reporting – The UPF sends the usage report of this specific URR when this trigger is set. In addition, the usage report is sent to any of the URRs linked to UPF when this trigger is set. The UPF sends the linked URR-Id along with the linked URR-Id IE.
- Packet Forwarding Control Protocol (PFCP) Session Deletion – A usage report generates (in a PFCP Deletion Response) for a URR due to the termination of the PFCP session. Similarly, a usage report generates (in a PFCP modification response) for a URR due to the removal of a specific URR.
- Update URR – This trigger is set when update URR request is received.
- Recalculate Measurement – The custom Information Element (IE) called Recalculate Measurement is added to the Update-URR process (URR-ID: Gz-Bearer) to support the PGWCDR generation due to the Max-LOSDV condition. After receiving the IE, UPF reconcile the Volume and Duration of the URR according to the URR-ID present inside the IE.

Note

- The total usage for the given URR-ID is reset, and if there are any linked URRs, the usage is recalculated.
- The timestamps for first-pkt-time, last-pkt-time and duration are also adjusted accordingly.
- The threshold values remain unchanged.

- Recalculate Interface - The custom IE called ‘Recalculate Interface’ is sent by SMF via N4 Modify Request. When Recalculate Interface IE is received along with Query URR(s) or Update URR(s) IEs in the same Sx Modify Request, Recalculate Interface IE will be processed independently in the end. After receiving the IE, the UPF iterates through all the URRs, fetching only the URRs mentioned as per the flags mentioned by SMF and execute the recalculate functionality. The UPF subsequently reconciles the Volume and Duration of all the retrieved URRs.

Note

- The usage accumulated for the URRs is discarded, and if there are any linked URRs, their usage is reconciled.
- The timestamps for first-pkt-time, last-pkt-time and duration are adjusted accordingly.
- The threshold values for the linked URRs remain unchanged.

IEs Supported for Online Charging Reporting

The following IEs support online charging:

- Volume Quota – The SMF requests the UPF to stop forwarding packets or allow forwarding some limited user plane traffic (based on the operator policy in UPF) with this IE. If no Volume Threshold is provisioned – to generate a usage report – and when the measured traffic reaches the quota, this IE is used.
- Time Quota - The SMF requests the UPF to stop forwarding packets or allow forwarding some limited user plane traffic (based on operator policy in UPF) with this IE. If no Volume Threshold is provisioned – to generate a usage report – and when the measured traffic reaches the quota, this IE is used.
- Monitoring Time – This IE is used by the SMF to send the time (UTC format) at which the UPF can re-apply the volume or time threshold. Also, the SMF sends any one of the Subsequent Volume, Time, Volume Quota, Time Quota, and Quota IEs, which is re-applied at the Monitoring Timestamp.
- FAR (Forwarding Action Rule) ID for Quota Action – This IE is used by the SMF to identify the substitute FAR the UPF applies – for the traffic that is associated to the URR – when any of the Volume, Time or quota is exhausted. This FAR requires the UPF to drop the packets or redirect the traffic toward a redirect destination.
- Subsequent Volume Threshold – When volume-based measurement is used and Monitoring Time IE is available, this IE is also present. The presence of this IE indicates the existence of the traffic volume value (the network resources usage reported by the UPF to the SMF) for this specific URR and the period after the Monitoring Time.
- Subsequent Time Threshold - When time-based measurement is used and Monitoring Time IE is available, this IE is also present. The presence of this IE indicates the existence of the time usage (the network resources usage reported by the UP function to the CP function) for this specific URR and the period after the Monitoring Time.
- Linked URR ID – When the linked usage reporting is required, this IE is used. It is possible to link multiple URR-IDs with an URR. Also, linked usage reporting is also sent in the Reporting Trigger IE.
- Measurement Method – The SMF specifies the measurement method of the network usage with the presence of this IE. The measurement method is based on volume and duration.
- Measurement Period – This IE is present to modify the measurement period.
- Periodic Reporting - When this trigger is set, the UPF sends resource usage report periodically to the SMF. The intervals that are required for periodic reporting are sent with the measurement period IE. When the trigger is set to 1, a request for periodic reporting is sent.
- Volume Threshold – This trigger is set when volume-based measurement is required. The SMF sends the traffic volume value along with the volume threshold IE, while the UPF sends the traffic usage report when the traffic volume is reached for the specific Usage Reporting Rule (URR). When the trigger is set to 1, a request for reporting – when the data volume usage reaches a volume threshold – is sent.
- Time Threshold - This trigger is set when time-based measurement is set. The SMF sends the time threshold value along with the time threshold IE, while the UPF sends resource usage report when the time threshold is reached for the specific URR. When the trigger is set to 1, a request for reporting – when the time usage reaches a time threshold - is sent.
- Start of Traffic – The UPF sends the Usage Report once the traffic starts for an application, when this trigger is set.
- Linked Usage Reporting - The UPF sends the usage report of this specific URR when this trigger is set. In addition, the usage report is sent to any of the URRs linked to UPF when this trigger is set. The UPF sends the linked URR-Id along with the linked URR-Id IE. When the trigger is set to 1, a request for linked usage reporting is sent.

Usage Reporting in PFCP Modification Response

The UPF sends session modification response after receiving session modification request based on the IEs received in the request message. The UPF includes usage report IE in the session modification response for the following scenarios:

- Query URR Handling—The URR-Id IE is included when the SMF requests immediate usage reports from the UPF in the session modification response (for the URR-Id present in this specific IE).
- Query All URRs (QAURR) Handling—The UPF sends the usage report with session modification response for all the URRs provisioned prior by the SMF for this PFCP session once it receives the QUARR flag set in PFCPSMReq-Flags IE from SMF.
- Update URR—The SMF updates the new value of the existing IE with the old value during the session modification procedure.
- Remove URR—During the session modification procedure, the SMF removes the IE, which is not received but was available earlier.

Usage Reporting for Online and Offline Charging

Usage Reporting for Online and Offline Charging is supported in the following ways:

- URR for online charging based on Rating-Group level even if the Service ID is present under Charging-Action. This behavior is seen when diameter ignore-service-id is configured under Credit Control Group.
- URR for offline charging based on a combination of Rating-Group level and Service ID, for static and predefined rules, as configured in the Charging-Action.

Both URRs are linked by the SMF. These URRs are linked such that when an online URR is reported, an offline URR is also reported.

Note If the Update URR does not contain any time or volume threshold values, that threshold is not monitored after that point.

Usage Reporting with Rating-Group and Service ID

The functionality enables usage reporting to the SMF with the Rating-Group (RG) and/or Service ID (SI) populated in the Usage Report IE within the Session Report Request.

The RG and SI are populated using proprietary PFCP IEs and are applicable for usage reporting of URRs associated only with Static and Predefined configured rules. The values are derived from the configured charging-action associated with the ruledefs, resulting in creation of the URRs during predefined activation or traffic hit for static rules.

Any change in RG/SI properties of the charging-action is reflected only in new URRs. The existing URRs associated with such charging-actions continue to report usage with the earlier RG+SI values.

UPF does not differentiate between usage reporting for Online and Offline URRs, and reports the RG+SI/RG/SI values configured in the charging-action, resulting in creation of the URRs.

NOTE: To know how SMF handles this functionality, refer *Dynamic Configuration Change Support* section in the *SMF Charging* chapter of *UCC 5G SMF Configuration and Administration Guide*.

Implementing the QAURR Flag

The SMF sets the QAURR flag of PFCPSMReq-Flags IE to request immediate usage reports for all the URRs previously provisioned earlier. Alternatively, SMF queries report for selected URR by sending URR-ID with Query URR IE. The UPF sends the usage report IE for corresponding URR with PFCP session modification response when the SMF sends the QAURR flag set in PFCPSMReq-Flag IE or URR-Id with Query URR IE.

Supported Functionality and Limitations

Basic call flow with Volume-Quota mechanism is supported with the following limitations:

- Dynamic Rules with Online Enabled is supported; both at Session-Setup and Mid-Session.
- Predefined Rules (dynamic-only) is supported; both at Session-Setup and Mid-Session. No restriction on configuring the "preemptively request".
- Static-rules with Online Charging are supported.
- Ignore-service-id is supported.
- Volume-Quota/Volume-Threshold mechanisms are supported.
- Event-Triggers (through which the Query URR occurs), and sending of usage information to the OCS is supported.
- The "updateURR" procedure, through the Sx/N4-Session-Modification procedure where the OCS grants a fresh Quota, is supported.
- Pending-Traffic-Treatment (PTT) Drop/Pass is supported with following limitations:
  - The scenarios currently supported are no-quota and quota-exhausted.
  - The trigger or re-authorization scenarios are not supported.
  - The PTT action (Forward/Drop) is considered after the quota-get is exhausted.
- Wall-Clock time-quota mechanism is supported.
- Other Time Quota Mechanisms (Discrete Time Period and Continuous-Time-Period) are not supported.
- Final-Unit-Indication Terminate mechanism is supported.
- FUI-Restrict is not supported.
- Server-Unreachable (SU) mechanism is now supported with minor change in behavior compared to non-5G UPF P-GW.
  - When an URR needs quota at UPF, the usage-report is generated to SMF and until the SMF responds with the linked SU_URR, the packets matching this URR are treated with Pending-Traffic-Treatment configuration.
  - When the SU Time Quota is used and it's reported to SMF for the Quota Exhaust, and if the session goes into Server-Unreachable state again, the time elapsed from the last Usage-Report is accounted in the usage.

- Pending-Traffic-Treatment Buffer mechanism is not supported.
- Quota-Hold-Time is supported.
- Quota-Consumption-Time mechanism is not supported.
- Quota-Validity-Time is supported.
- Configuring different "rating-group" value other than the "content-id" is supported.
  - The RG 0 is not supported.
- Trigger to PCF for the Out-of-Credit, Reallocation-of-Credit events are not qualified.



Important Event-trigger Out-of-Credit toward PCF is validated with a limitation of having only one time Grant-Quota (Keeping Total Volume and Granted Volume at same value).

- Service-Specific-Units are not supported.
- Tariff-Time change is supported as per 3GPP specification.
- FUI-Redirect is supported with following limitations:
  - Redirection for HTTPs is not supported.
  - The FUI-Redirect with Filter-IDs/Filter-Rules are not supported.
  - The WSP Protocol is not supported.
  - The **redirect-require-user-agent** CLI command is not supported; the redirection continues to work even if the user-agent is not present.
  - Appending the original URL is not supported.
  - Token-based mechanism, to come out of Redirection, is not supported. To end the redirection in 5G UPF, OCS sends Redirect Validity-Time or RAR.
  - FUI-Redirection is supported only for the URL, similar to the behavior in non-5G UPF architecture.
  - Check pointing of FUI Redirection URL is not supported.

PTT no-quota Limited Pass

This feature allows the subscriber to use the network while waiting for the response from OCS. The Limited-Pass configuration allows to specify the Volume which the subscriber can consume while waiting for the quota-response from OCS. The usage is accounted in the respective charging bucket and is adjusted against the next-quota allocation.

Use the following CLI commands to enable the feature:

```
configure
  active-charging service service_name
    credit-control
```

```
pending-traffic-treatment noquota limited-pass volume volume
end
```

Limited Pass Volume is used only for **noquota** case (Rating Group (RG) seeking quota for the first time) and not for **quota-exhausted**. Limited Pass Volume isn't used for subsequent credit requests.

The traffic is allowed to pass until the Limited-Pass Volume gets exhausted. The usage is counted in the respected charging-bucket and adjusted against the "Quota" granted. If the "Quota" allocation is less than the actual usage, immediate reporting towards OCS with the usage-report occurs requesting for more quota allocation. The subsequent incoming packets are handled as per the "quota-exhausted" PTT configuration.

If the Limited Pass Volume is NOT exhausted before the OCS responds with denial of quota, traffic is blocked after the OCS response. The gateway reports usage on Limited-Pass Volume in next SX_SESSION_REPORT_REQUEST.

If the Limited Pass Volume is exhausted before the OCS responds, then the subsequent incoming packets for the session are dropped until quota is granted from OCS.

The default pending-traffic-treatment for **noquota** is Drop. The **default pending-traffic-treatment noquota** command removes any Limited Pass Volume size configured.

PTT quota exhaust Limited Pass

Quota Exhausted Limited pass is proposed as an alternative to the Quota Exhausted Buffer due to the practical issues of the latter in the high-speed network. Buffering requires packet buffering for large number of packets at the gateway. The large number of packets can result in risking to run out of memory affecting the bandwidth speed. So, Limited Pass is an alternate to the Buffer option. Limited Pass allows the traffic to pass through until the configured limit on the Quota-Exhaust scenarios.

Use the following CLI command to enable the feature:

```
configure
  active-charging service service_name
    credit-control
     pending-traffic-treatment quota-exhausted limited-pass volumevolume

     end
```

Note The above CLI is only applicable for the 5G UPF architecture.

After the Limited-Pass volume exhausts, the further packets drop until the quota allocation comes.

Limited Pass allows the traffic until the Limited-Pass volume exhausts. The Limited Pass counts and adjusts the usage in the respective charging-bucket against the "Quota" granted. If the "Quota" allocation is less than the actual usage, there’s immediate reporting toward OCS with the usage-report and asking for more quota allocation.

If the limited pass volume doesn’t exhaust before the OCS responds with denial of quota, there’s traffic blockage after the OCS response. Gateway reports the usage in SX_SESSION_REPORT_REQUEST.

If the limited pass volume exhausts before the OCS responds, then further incoming packets for the session drop until grant quota from OCS.

The default pending-traffic-treatment for quota-exhausted is drop. The default pending-traffic-treatment quota-exhausted command removes any Limited Pass Volume size configured.

Tariff Time Support

The Tariff switch time functionality is applied when a subscriber switch from one tariff plan to another.

The Tariff-Time-Change AVP is used to determine the tariff switch time, and the Monitoring-Time IE is used to support the Tariff Time support functionality.

After a tariff timer expiry, the Gateway accumulates the usage separately in a charging bucket and continues to consume from the original quota value. At the time of next reporting (Quota exhausted or another control events), the Gateway reports both usages (before and after tariff time change) for the same Charging Bucket.

The first reporting of this charging-bucket will have the Reporting-Reason: Monitoring Time and the second bucket will contain the last reporting reason, and the quota usage after the tariff-timer expiry.

The data traffic usage can be split into resource usage before a tariff switch and resources used after a tariff switch. The Tariff-Change-Usage AVP is used within the Used-Service-Units AVP to distinguish reported usage before and after the tariff time change.

Limitations

Following are the known limitations of this feature:

- Only one tariff time per RG/Service ID combination is supported.
- Allocation of different quota before and after tariff time change isn't supported. This functionality isn’t in compliance with the 3GPP standards.

TCP Maximum Segment Size

TCP/IP Stack always inserts Maximum Segment Size (MSS) field in the header. This causes difference in MSS insertion behavior with and without TCP Proxy.

Using **tcp mss** configurations, TCP MSS can be limited if already present in the TCP SYN packets. If there are no errors detected in IP header or TCP mandatory header, and there are no memory allocation failures, TCP optional header is parsed. If TCP MSS is present in the optional header and its value is greater than the configured MSS value, the value present in the TCP packet is replaced with the one that is configured.

If the TCP optional header is not present in the SYN packet and there are no errors in already-present TCP header, the configured TCP MSS value is inserted while sending out the current packet.

Configuring Credit Control for Usage Reporting

This configuration enables to accept/ignore service ID in the Service-Identifier AVP defined in the Diameter dictionaries.

```
configure
   require active charging
   active-charging service service_name
      credit-control group group_name
         diameter ignore-service-id
         end
```

- **diameter ignore-service-id** : This command can be used to disable the usage of the Service-Identifier AVP for Gy interface implementations even if any of the Diameter dictionaries support the

Service-Identifier AVP, and if this AVP should not be used for Gy interactions but must be present in GCDRs/eGCDRs.

Configuring ACS Rulebase for Usage Reporting

This section describes how to create, configure, or delete an ACS rulebase. A rulebase is a collection of protocol rules to match a flow and associated actions to be taken for matching flow. The default rulebase is used when a subscriber/APN is not configured with a specific rulebase to use.

Rulebase configuration is the one that combines all the specified configurations together to construct the static and predefined PCC rules.

```
configure
   active-charging service service_name
      rulebase rulebase_name
         action priority action_priority { [ dynamic-only ] |
static-and-dynamic | timedef timedef_name ] { group-of-ruledefs
ruledefs_group_name | ruledef ruledef_name } charging-action charging_action_name [
 monitoring-key monitoring_key ] [ description description ] }
         cca quota { holding-time holding_time content-id content_id |
retry-time retry_time [ max-retries retries ] }
         credit-control-group cc_group_name
         dynamic-rule order { always-first | first-if-tied }
         egcdr threshold { interval interval [ regardless-of-other-triggers
 ] | volume { downlink | total | uplink } bytes }
         route priority route_priority ruledef ruledef_name analyzer { dns |
file-transfer | ftp-control | ftp-data | h323 | http | imap | mipv6 | mms
 | pop3 | pptp | radius | rtcp | rtp | rtsp | sdp | secure-http | sip [
advanced | basic-and-advanced ] | smtp | tftp | wsp-connection-less |
wsp-connection-oriented } [ description description ]
         tcp check-window-size
         tcp mss tcp_mss { add-if-not-present | limit-if-present |
limit-if-present add-if-not-present }
         tcp packets-out-of-order { timeout timeout_duration | transmit [
after-reordering | immediately ] }
         end
```

NOTES:

- **rulebase** *rulebase_name*: Specifies the name of the ACS rulebase. *rulebase_name* must be an alphanumeric string of 1 to 63 characters.
- **action priority** *action_priority* **{ [dynamic-only] | static-and-dynamic | timedef** *timedef_name* **] { group-of-ruledefs** *ruledefs_group_name* **| ruledef** *ruledef_name* **} charging-action** *charging_action_name* **[monitoring-key** *monitoring_key* **] [description** *description* **] }**: Configures the priority order in which ruledefs are matched and the associated charging action.
  - *priority* must be an integer value in the range of 1-65535.
  - *monitoring_key* must be an integer value in the range of 100000-4000000000.

- **cca quota { holding-time** *holding_time* **content-id** *content_id* **| retry-time** *retry_time* **[max-retries** *retries* **] }**: Configures the quota for the online charging.
  - *holding_time*: must be an integer value in the range of 1-4000000000
  - *content_id*: must be an integer value in the range of 1-2147483647
  - *retry_time*: must be an integer value in the range of 0-86400
  - *retries*: must be an integer value in the range of 1-65535
- **cca quota time-duration algorithm { consumed-time** *seconds* **[plus-idle] | continuous-time-periods** *seconds* **| parking-meter** *seconds* **} [content-id** *content_id* **]**
  - consumed-time: must be an integer value in the range of 1-4294967295
  - content-id: must be an integer value in the range of 1-2147483647
  - continuous-time-periods: must be an integer value in the range of 1-4294967295
  - parking-meter: must be an integer value in the range of 1-4294967295
- **credit-control-group** *cc_group_name*: Configures the online charging parameters used by this rulebase. *cc_group_name* must be an alphanumeric string of 1 to 63 characters.
- **dynamic-rule order**: Configures the order of dynamic rule matching vs the static rules in a rulebase.
- **egcdr threshold { interval** *interval* **[regardless-of-other-triggers] | volume { downlink | total | uplink } bytes }**: Configures the threshold for offline charging.
  - **interval**: must be an integer value in the range of 60-40000000.
  - **downlink**: must be an integer value in the range of 100000-4000000000. Default: 4000000000.
  - **uplink**: must be an integer value in the range of 100000-4000000000. Default: 4000000000.
  - **total**: must be an integer value in the range of 100000-4000000000.
- **route priority** *route_priority* **ruledef** *ruledef_name* **analyzer { dns | file-transfer | ftp-control | ftp-data | h323 | http | imap | mipv6 | mms | pop3 | pptp | radius | rtcp | rtp | rtsp | sdp | secure-http | sip [advanced | basic-and-advanced] | smtp | tftp | wsp-connection-less | wsp-connection-oriented } [description** *description* **]**: This command is used only on UPF.
  - *route_priority* must be an integer value in the range of 0-65535.
  - *ruledef_name* must be an alphanumeric string of 1 to 63 characters.
- **tcp check-window-size**: This command is used only on UPF.
- **tcp mss** *tcp_mss*: This command is used only on UPF. *tcp_mss* must be an integer value in the range of 496-65535.
  - **add-if-not-present** : Specifies to add the TCP MSS if not present in the packet.
  - **limit-if-present** : Specifies to limit the TCP MSS if present in the packet.
  - **limit-if-present add-if-not-present** : Specifies to limit the TCP MSS if present, else, adds it to the packets.

Note The **tcp mss** *tcp_mss* **limit-if-present add-if-not-present** CLI command is available in 2021.02.0 and later releases.

- **tcp packets-out-of-order { timeout** *timeout_duration* **| transmit [after-reordering | immediately] }**: This command is used only on UPF.
  - *timeout_duration* must be an integer value in the range of 100-30000. Default value is 5000.

Sample Configuration

```
active-charging service acs
    ruledef ip-any-rule
      ip any-match = TRUE
    #exit
    urr-list upf
      rating-group 10 ser 10 urr-id 10
      rating-group 10 urr-id 50
    #exit
    charging-action starent
      content-id 10
      service-identifier 10
      billing-action egcdr
      cca charging credit rating-group 10
    exit
    credit-control group CCG
      diameter ignore-service-id
    #exit
    rulebase starent
      billing-records egcdr
      action priority 30 ruledef ip-any-rule charging-action starent
      egcdr threshold interval 3600
      egcdr threshold volume total 200000
      egcdr threshold volume downlink 100000 uplink 100000
      dynamic-rule order first-if-tied
      credit-control-group CCG
    #exit
#exit

context ISP
    apn starent.com
      accounting-mode gtpp
      gtpp group my_grp accounting-context ISP
      ip context-name ISP
    #exit
    gtpp group my_grp
      gtpp egcdr service-data-flow threshold interval 1200
      gtpp egcdr service-data-flow threshold volume downlink 13000
      gtpp egcdr service-data-flow threshold volume uplink 17000
      gtpp egcdr service-data-flow threshold volume total 22222
    #exit
end
```

Monitoring and Troubleshooting

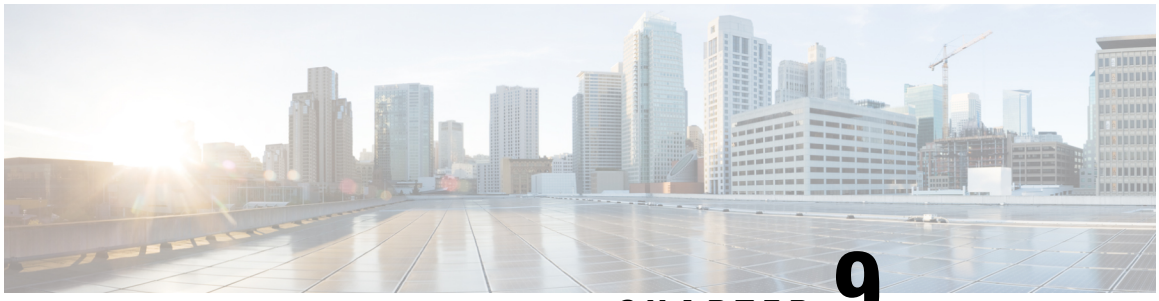
Show Commands and/or Outputs

This section provides information about the show CLI commands that are available in support of the feature.

show-user-plane-service statistics rulebase name <name>

Use this CLI command to see the following fields that are available in support of TCP Maximum Segment Size (MSS) feature:

- TCP MSS Inserted Pkts: Displays the total number of MSS Inserted packets.
- TCP MSS Limited Pkts: Displays the total number of TCP MSS Limited packets.



CHAPTER 9

Cisco Ultra Traffic Optimization with VPP

- Feature Summary and Revision History, on page 81
- Feature Description, on page 82
- RCM Support, on page 82
- Sending the GBR or MBR Values to Cisco Ultra Traffic Optimization , on page 83
- How it Works, on page 83
- Show Commands and Outputs, on page 85
- Sample Configuration, on page 91

Feature Summary and Revision History

Summary Data

Table 18: Summary Data

Applicable Product(s) or Functional Area	5G-UPF
Applicable Platform(s)	VPC-SI
Feature Default Setting	Disabled - Configuration Required
Related Changes in this Release	Not Applicable
Related Documentation	Not Applicable

Revision History

Table 19: Revision History

Revision Details	Release
First introduced.	2022.01.1

Feature Description

The UPF supports Cisco Ultra Traffic Optimization (CUTO) on Vector Packet Processing (VPP).

The Cisco Ultra Traffic Optimization is a RAN optimization technology that increases the subscriber connection speeds in congested cells and, as a result, increases the cell capacity significantly. The result is an optimized RAN, where Mobile Network Operators (MNOs) can deploy fewer cells, on an ongoing basis, and absorb more traffic growth while meeting network quality targets.

Large traffic flows, such as Adaptive Bit Rate (ABR) video, saturate radio resources and swamp the eNodeB scheduler. The Cisco Ultra Traffic Optimization employs machine learning algorithms to detect large traffic flows (such as video) in the network. It also optimizes the Delivery of those flows to mitigate the network congestion without changing the user quality (that is, video works the same for you). In other words, by employing software intelligence at the network core, Cisco Ultra Traffic Optimization mitigates the overwhelming impact the video has on the RAN.

The resulting benefits are seen in congested network sites. The Cisco Ultra Traffic Optimization:

- Increases average user throughput.
- Increases congested cell site capacity.
- Reduces scheduler latency.
- Maintains user quality of experience even when more users and more traffic share a cell.
- Is measured directly by eNodeB performance counters (for example, average UE throughput, scheduler latency). These are the key performance indicators that are used for network capacity planning.
- Provides permanent savings in RAN investment requirements.
- Is integrated in the Cisco StarOS P-GW.
- Requires no new hardware or cabling complexity - it can be turned on for a market in an hour.
- Supports HTTP or HTTPS, and QUIC traffic.

Licensing

The Cisco Ultra Traffic Optimization with VPP is a licensed Cisco solution. Contact your Cisco account representative for detailed information on specific licensing requirements. For information on installing and verifying licenses, refer to the *Managing License Keys* section of the *Software Management Operations* chapter in the *VPC-SI System Administration Guide*.

RCM Support

This feature enables the Redundancy and Configuration Management (RCM) support for the Cisco Ultra Traffic Optimization (CUTO). All relevant configuration to enable CUTO using service scheme and application of the CUTO profile or policy on UPF is supported using RCM.

Sending the GBR or MBR Values to Cisco Ultra Traffic Optimization

If the flow level MBR is greater than the APN-AMBR for a non GBR bearer, traffic is throttled at APN-AMBR. In such a case APN-AMBR is sent as the upper limit to the CUTO library. If there is no valid flow level MBR specific to the flow, APN-AMBR is sent as the upper limit to the CUTO library.

For a GBR bearer, the flow level GBR is sent as the lower limit and flow level MBR is sent as the upper limit to the CUTO library.

Cisco Ultra Traffic Optimization Library Deinitialization

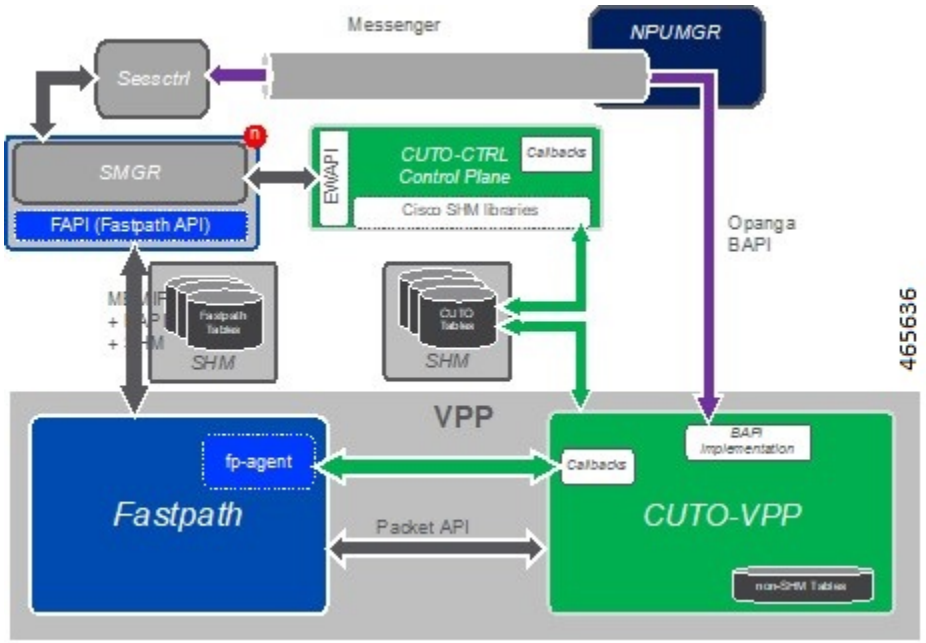
This feature currently doesn't support the Deinitialization. Deinitialization happens when the Cisco Ultra Traffic Optimization (CUTO) license is removed from the system.

How it Works

Architecture

The following figure illustrates the architecture of Cisco Ultra Traffic Optimization on VPP.

Figure 10: Architecture



CUTO-CTRL

- CUTO-CTRL receives guidance and requests from SMGR through the East-West API (EWAPI), through which clients (SMGR instances) are registered and deregistered, and new streams or flows are created and terminated.
- CUTO-CTRL manages a set of shared memory (SHM) tables using a North-South API (NSAPI) consisting of Cisco-provided SHM infrastructure.
- It is through this SHM environment that CUTO-VPP can read and write content that is visible to both CUTO-VPP and CUTO-CTRL.
- The SHM is used for all high volume, scalable/mutable content necessary for the high-performance configuration and administration of the CUTO solution in VPP.

NPUMGR

NPUMGR is the management layer that is responsible for the overall VPP operation. It sends Binary API (BAPI) requests to CUTO-VPP for initialization, global runtime configuration, and policy configuration.

SMGR

SMGR is the main subscriber control plane. There are N SMGR instances, and all instances are managed by the SessCtrl. In the context of VPPMOB/Fastpath, the SMGR instances are also known as “Clients”, and each client has a unique ID.

SMGR issues policy directives to SessCtrl through the Messenger tunnel, and sends updated Policy guidance to CUTO-VPP through the Binary API.

SMGR communicates with Fastpath for pre-existing functionality with a set of MEMIF, Binary API, and shared memory (SHM) infrastructures.

Session Control (SessCtrl)

Session Control is the management layer responsible for overseeing the set of SMGR instances.

The BAPI requests are tunnelled from SessCtrl to NPUMGR through Messenger.

Fastpath

VPP is responsible for packet processing. Fastpath performs subscriber-related packet processing within the VPP environment. Subscriber flows are divided into unidirectional Streams, and a Stream conduit is the pipeline of functions through which a packet is transformed and egressed from subscriber processing.

A packet API between the Fastpath and CUTO-VPP facilitates the exchange of packets traversing the Fastpath conduit.

CUTO-VPP

- CUTO-VPP is the packet processing engine in the UPF.
- In fastpath, Cisco Ultra Traffic Optimization is applied to packets on a stream configured with its operation.
- Packets are sent from the Stream conduit to a particular CUTO-VPP operation, and after some potential delay (0-N milliseconds), traffic is returned to the same Conduit.
- Packets are never dropped by the Cisco Ultra Traffic optimization library.

CUTO-TODR

Traffic Optimization Data Records (TODR) can only be generated as events, and are enabled only when the configuration is available.

Limitations

The Cisco Ultra Traffic Optimization feature has the following limitations:

- CUTO configuration changes done in Service Schema do not take effect immediately for existing flows.
- Cisco Ultra Traffic Optimization VPP global deinitialization is not supported.
- Bearer-related triggers for enabling Cisco Ultra Traffic Optimization are not supported.
- Rule match change trigger must be configured for CUTO in UPF.
- Enabling/Disabling of Traffic optimization is not supported on "loc-update" trigger.
- Removal of CUTO license doesn't trigger global deinitialization. CUTO configurations must be removed to disengage CUTO functionality for new flows.

Show Commands and Outputs

This section provides information regarding show commands and their outputs in support of Cisco Ultra Traffic Optimization.

For information on other supporting show commands, refer to *Monitoring and Troubleshooting* section under the *Cisco Ultra Traffic Optimization* chapter in the *P-GW Administration Guide*.

Show Commands and Outputs

show user-plane-service traffic-optimization counters sessmgr all

The output of this command includes the following fields:

TCP Traffic Optimization Flows:

- Active Normal Flow Count
- Active Large Flow Count
- Active Managed Large Flow Count
- Active Unmanaged Large Flow Count
- Total Normal Flow Count
- Total Large Flow Count
- Total Managed Large Flow Count
- Total Unmanaged Large Flow Count
- Total IO Bytes

- Total Large Flow Bytes
- Total Recovered Capacity Bytes
- Total Recovered Capacity ms

UDP Traffic Optimization Flows:

- Active Normal Flow Count
- Active Large Flow Count
- Active Managed Large Flow Count
- Active Unmanaged Large Flow Count
- Total Normal Flow Count
- Total Large Flow Count
- Total Managed Large Flow Count
- Total Unmanaged Large Flow Count
- Total IO Bytes
- Total Large Flow Bytes
- Total Recovered Capacity Bytes
- Total Recovered Capacity ms

show user-plane-service traffic-optimization info

The output of this command includes the following fields:

- CUTO Ctrl Library Version
- CUTO VPP Library Version
- Mode
- Configuration
  - Data Records (TODR)
  - Statistics Options
  - EFD Flow Cleanup Interval
  - Statistics Interval

show user-plane-service traffic-optimization policy all

The output of this command includes the following fields:

- Policy Name
- Policy-Id

- Bandwidth-Mgmt
  - Backoff-Profile
  - Min-Effective-Rate
  - Min-Flow-Control-Rate
- Curbing-Control:
  - Time
  - Rate
  - Max-Phases
  - Threshold-Rate
- Heavy-Session:
  - Threshold
  - Standard-Flow-Timeout
  - Seed-Time
- Detection-Mode
- Link-Profile:
  - Initial-Rate
  - Max-Rate
  - Peak-Lock
- Session-Params:
  - Tcp-Ramp-Up
  - Udp-Ramp-Up
- Total traffic-optimization-policies found

show user-plane-service traffic-optimization init-status sessmgr all

The output of this command includes the following fields:

- Sessmgr Instance
  - CUTO Ctrl Library
  - CUTO VPP Library

show user-plane-service traffic-optimization statistics sessmgr all

The output of this command includes the following fields:

- Traffic Optimization Stats:
  - Uplink Drop
  - Uplink Hold
  - Uplink Forward
  - Uplink Rx Pkts
  - Uplink Tx Pkts
  - Dnlink Drop
  - Dnlink Hold
  - Dnlink Forward
  - Downlink Rx Pkts
  - Downlink Tx Pkts
  - TODRs Generated
- CUTO Control Plane Stats:
  - Init:
    - Success:
      - SESSMGR-CP request
      - SESSMGR-CP response
      - CP-VAPI request
      - CP-VAPI response
    - Failure:
      - SESSMGR-CP request
      - SESSMGR-CP response
      - CP-VAPI request
      - CP-VAPI response
      - CP-VAPI response error
  - Deinit:
    - Success:
      - SESSMGR-CP request
      - SESSMGR-CP response
    - Failure:
      - SESSMGR-CP request

- SESSMGR-CP response

- Cumulative:
  - Success:
    - Total TCP create stream request
    - Total TCP update stream request
    - Total TCP delete stream request
    - Total UDP(QUIC) create stream request
    - Total UDP(QUIC) create stream request
    - Total UDP(QUIC) delete stream request
    - Total format TODR request
    - Total delete 3pid notify request
  - Failure:
    - Total TCP create stream request
    - Total TCP update stream request
    - Total TCP delete stream request
    - Total UDP(QUIC) create stream request
    - Total UDP(QUIC) create stream request
    - Total UDP(QUIC) delete stream request
    - Total format TODR request
    - Total delete 3pid notify request

- Current
  - Total Active streams
  - Active TCP streams
  - Active UDP(QUIC) streams

Bulkstats

The following existing bulk statistics are supported by Cisco Ultra Traffic Optimization in UPF:

Bulk Statistics	Description
cuto-uplink-drop	Indicates the total number of uplink packets dropped by CUTO library

Bulk Statistics	Description
cuto-uplink-hold	Indicates the total number of uplink packets held by CUTO library
cuto-uplink-forward	Indicates the total number of uplink packets forwarded by CUTO library
cuto-uplink-rx	Indicates the total number of uplink packets received by CUTO library
cuto-uplink-tx	Indicates the total number of uplink packets sent by CUTO library
cuto-dnlink-drop	Indicates the total number of downlink packets dropped by CUTO library
cuto-dnlink-hold	Indicates the total number of downlink packets held by CUTO library
cuto-dnlink-forward	Indicates the total number of downlink packets forwarded by CUTO library
cuto-dnlink-rx	Indicates the total number of downlink packets received by CUTO library
cuto-dnlink-tx	Indicates the total number of downlink packets sent by CUTO library
cuto-todrs-generated	Indicates the total number of TODRs generated.
tcp-active-normal-flow-count	Indicates the number of TCP active-normal-flow count for Cisco Ultra Traffic Optimization.
tcp-active-large-flow-count	Indicates the number of TCP active-large-flow count for Cisco Ultra Traffic Optimization.
tcp-active-managed-large-flow-count	Indicates the number of TCP active-managed-large-flow count for Cisco Ultra Traffic Optimization.
tcp-active-unmanaged-large-flow-count	Indicates the number of TCP active-unmanaged-large-flow count for Cisco Ultra Traffic Optimization.
tcp-total-normal-flow-count	Indicates the number of TCP total-normal-flow count for Cisco Ultra Traffic Optimization.
tcp-total-large-flow-count	Indicates the number of TCP total-large-flow count for Cisco Ultra Traffic Optimization.
tcp-total-managed-large-flow-count	Indicates the number of TCP total-managed-large-flow count for Cisco Ultra Traffic Optimization.
tcp-total-unmanaged-large-flow-count	Indicates the number of TCP total-unmanaged-large-flow count for Cisco Ultra Traffic Optimization.
tcp-total-io-bytes	Indicates the number of TCP total-IO bytes for Cisco Ultra Traffic Optimization.

Bulk Statistics	Description
tcp-total-large-flow-bytes	Indicates the number of TCP total-large-flow bytes for Cisco Ultra Traffic Optimization.
tcp-total-recovered-capacity-bytes	Indicates the number of TCP total-recovered capacity bytes for Cisco Ultra Traffic Optimization.
tcp-total-recovered-capacity-ms	Indicates the number of TCP total-recovered capacity ms for Cisco Ultra Traffic Optimization.
udp-active-normal-flow-count	Indicates the number of UDP active-normal-flow count for Cisco Ultra Traffic Optimization.
udp-active-large-flow-count	Indicates the number of UDP active-large-flow count for Cisco Ultra Traffic Optimization.
udp-active-managed-large-flow-count	Indicates the number of UDP active-managed-large-flow count for Cisco Ultra Traffic Optimization.
udp-active-unmanaged-large-flow-count	Indicates the number of UDP active-unmanaged-large-flow count for Cisco Ultra Traffic Optimization.
udp-total-normal-flow-count	Indicates the number of UDP total-normal-flow count for Cisco Ultra Traffic Optimization.
udp-total-large-flow-count	Indicates the number of UDP total-large-flow count for Cisco Ultra Traffic Optimization.
udp-total-managed-large-flow-count	Indicates the number of UDP total-managed-large-flow count for Cisco Ultra Traffic Optimization.
udp-total-unmanaged-large-flow-count	Indicates the number of UDP total-unmanaged-large-flow count for Cisco Ultra Traffic Optimization.
udp-total-io-bytes	Indicates the number of UDP total-IO bytes for Cisco Ultra Traffic Optimization.
udp-total-large-flow-bytes	Indicates the number of UDP total-large-flow bytes for Cisco Ultra Traffic Optimization.
udp-total-recovered-capacity-bytes	Indicates the number of UDP total-recovered capacity bytes for Cisco Ultra Traffic Optimization.
udp-total-recovered-capacity-ms	Indicates the number of UDP total-recovered capacity ms for Cisco Ultra Traffic Optimization.

Sample Configuration

Sample configuration to enable the CUTO feature:

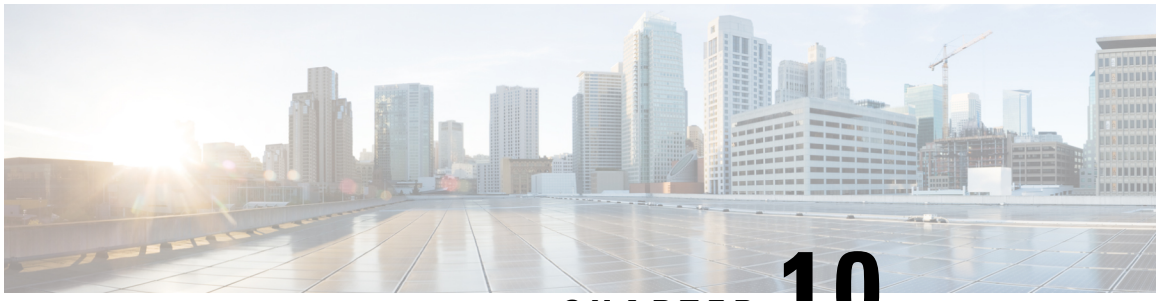
```
configure
  active-charging service ACS
```

```
    trigger-action TA1
       traffic-optimization policy custom1
    #exit
    trigger-condition TC1
      rule-name = dynamic-rule2
    #exit
    service-scheme SS1
      trigger rule-match-change
        priority 5 trigger-condition TC1 trigger-action TA1
      #exit
    subs-class SB1
      rulebase = cisco
    #exit
    subscriber-base default
      priority 5 subs-class SB1 bind service-scheme SS1
    #exit
    traffic-optimization-profile
      mode active
      data-record
    #exit
    traffic-optimization-policy custom1
      bandwidth-mgmt min-effective-rate 800 min-flow-control-rate 250
      heavy-session threshold 200000
      link-profile max-rate 20000
    #exit
    traffic-optimization-policy default
    #exit
end
```

CUTO-TODR

Sample configuration to enable the CUTO-TODR:

```
context ISP
edr-module active-charging-service
file name EDR directory TODR_CUTO rotation volume 51200 headers
cdr use-harddisk
```



CHAPTER 10

Collection and Reporting of Usage Data over N4 Interface

This chapter covers the following topics:

- Feature Summary and Revision History, on page 93
- Feature Description, on page 94
- How it Works, on page 94
- Configuration to Collect and Report Volume Measurement over N4 Interface, on page 95

Feature Summary and Revision History

Summary Data

Table 20: Summary Data

Applicable Product(s) or Functional Area	5G-UPF
Applicable Platform(s)	VPC-SI SMI
Feature Default Setting	Enabled – Always-on
Related Changes in this Release	Not Applicable
Related Documentation	Not Applicable

Revision History

Table 21: Revision History

Revision Details	Release
First Introduced	2020.02.0

Feature Description

With this release, the User Plane Function (UPF) supports offline charging and reporting of usage data over the N4 interface.

Here, the SMF controls the collection and reporting of usage data by creating necessary PDRs and URRs, and associates the URRs with its relevant PDRs defined for a PFCP session. It also controls data usage reporting at an IP-CAN bearer level, IP-CAN session, TDF session, SDF, or at an application level.

The URR consists of the usage measurement method, reporting triggers, threshold, and quota values.

Important In this release, only URR creation is supported during PFCP session establishment.

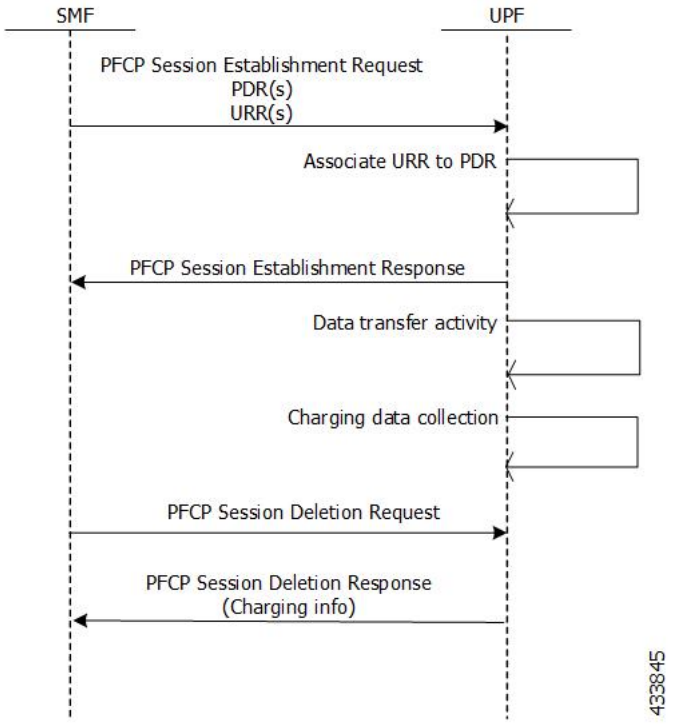
How it Works

This section describes how UPF supports offline charging of usage data.

To implement offline charging, the charging information is sent to the SMF only during PFCP session deletion.

Time and volume-based reporting is supported in the offline charging implementation. The following call flow illustrates offline charging in UPF.

Figure 11: Offline Charging in UPF



During the PFCP session deletion, UPF transfers the following charging information to the SMF:

- Timestamp of the first and last data packet
- Duration measurement – This IE specifies the time difference between URR creation and usage-reporting
- Volume measurement – This IE specifies the uplink data, downlink data and the total bytes transferred from the UPF to gNodeB.

Standards Compliance

UPF support for collection and reporting of data is compliant with the following standards:

- 3GPP TS 29.244 - LTE; Interface between the Control Plane and the User Plane of EPC Nodes
- 3GPP TS 23.501 - 5G; System Architecture for the 5G System
- 3GPP TS 23.502v - 5G; Procedures for the 5G System

Configuration to Collect and Report Volume Measurement over N4 Interface

This section describes the configuration required to collect and report volume measurement (usage data). However, to achieve this, SMF-based configurations for volume measurement needs to be configured.

The following SMF-based configuration is required to send volume measurement data in the URR by the UPF.

Configuring Charging Action for a Required Billing Action

Use the following configuration to configure charging-action for a required billing-action:

```
configure
   require active-charging
   active-charging service service_name
      charging-action charging_action_name
         billing-action interface_name
         end
```

NOTES:

- **billing-action**: Enables the specified billing type. The supported interfaces are:
  - **egcdr**: Enables the GGSN charging data record.
  - **radius**: Adds the data to the RADIUS accounting record.
  - **rf**: Enables Rf accounting.

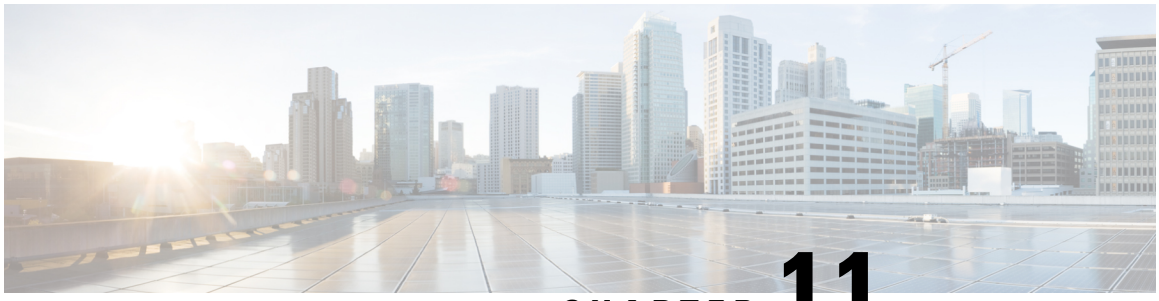
Associating a Charging Action with a Rulebase

Use the following configuration to associate a charging action with a rulebase:

```
configure
   require active-charging
   active-charging service service_name
      rulebase rulebase_name
         billing-records interface_name
         action priority priority_value ruledef ruledef_name charging-action
charging_action_name
         end
```

NOTES:

- **rulebase**: Enables the Active Charging Service Rulebase configuration.
- **billing-records**: Enables the generation of billing records. The supported interface is **egcdr**
- **action**: Decides the action to be taken on the ruledef.
- **priority**: Assigns priority to a ruledef in the rulebase. Priority must be a unique integer value ranging from 1 to 65535.
- **ruledef**: Specifies the ruledef.
- **charging-action**: Specifies the charging action.



CHAPTER 11

Control Plane-Initiated N4 Association Support

This chapter covers the following topics:

- Feature Summary and Revision History, on page 97
- Feature Description, on page 98
- How it Works, on page 98
- Configuring the CP-Initiated N4 Association Setup Feature, on page 98

Feature Summary and Revision History

Summary Data

Table 22: Summary Data

Applicable Product(s) or Functional Area	5G-UPF
Applicable Platform(s)	VPC-SI SMI
Feature Default Setting	Disabled – Configuration Required
Related Changes in this Release	Not Applicable
Related Documentation	Not Applicable

Revision History

Table 23: Revision History

Revision Details	Release
First Introduced.	2020.02.0

Feature Description

SMF initiated N4 Association Setup Procedure

The N4 association set up procedure sets up an N4 association between the Session Management Function (SMF) and User Plane Function (UPF). It enables the SMF to use the UPF resources to establish the N4 sessions. The SMF and UPF exchange the supported functionalities on each side during this procedure.

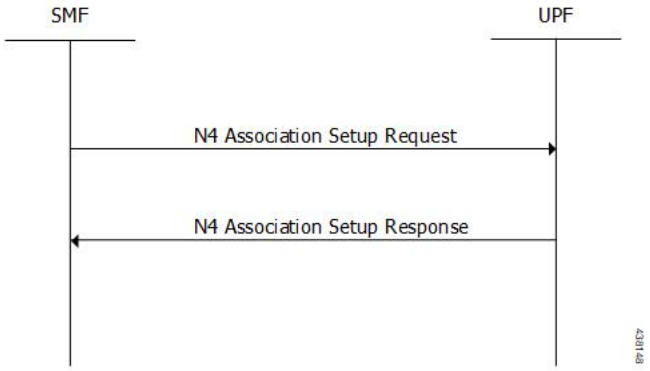
How it Works

The N4 association setup is initiated by the SMF. The setting of N4 association setup procedure is controlled through **sx-association initiated-by-cp** CLI command in the Control Plane Group Configuration mode. By default, the configuration is set to support the UPF-initiated N4 association setup procedure.

Call Flows

Session Management Function Initiated N4 Association Setup Procedure

The following call flow depicts the SMF-initiated N4 Association Setup procedure.



Step	Description
1	The SMF initiates the N4 Association Setup procedure to request the setup of an N4 association towards a UPF prior to establishing a first N4 session on this UPF.
2	After receiving an N4 Association Setup Request, the UPF sends an N4 Association Setup Response.

Configuring the CP-Initiated N4 Association Setup Feature

This section describes how to configure the CP-Initiated N4 Association Setup feature.

Configuring this feature involves using the "**sx-association initiated-by-cp**" CLI command in the Control Plane Group Configuration mode. The default configuration is UPF-initiated N4 association setup procedure.

Use the following configuration to configure the N4 association setup feature.

```
configure
   context
      control-plane-group group_name
         peer-node-id ipv4-address ip_address interface n4
         sx-association { initiated-by-cp | initiated-by-up }
         end
```

NOTES:

- **initiated-by-cp:** This keyword is used to initiate the Sx association request through control plane.
- **initiated-by-up:** This keyword is used to initiate the Sx association request through user plane.
- By default, the UPF-initiated N4 association setup procedure is configured.
- To revert to the default setting, use the **no sx-association** command.

CP-Initiated N4 Association Setup Feature OAM Support

This section describes operations, administration, and maintenance information for this feature.

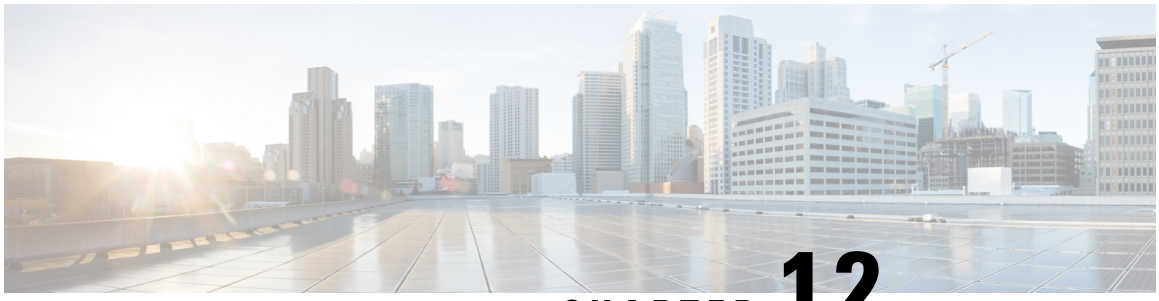
Show Command Support

Use the following show command to verify the CP-initiated N4 Association Setup feature configuration.

show control-plane-group all

The following is a sample output of the show command.

```
show control-plane-group all
Control Plane Group
---------------------------------
Name           : default
Sx-Association : initiated-by-up
Name           : default
Sx-Association : initiated-by-up
Node-Id        : 209.165.200.230
Interface      : N4
```



CHAPTER 12

Converged Datapath

- Feature Summary and Revision History, on page 101
- Feature Description, on page 102
- How it Works, on page 103
- Charging for Converged Calls, on page 116
- Configuring Converged Datapath, on page 118
- Monitoring and Troubleshooting, on page 119

Feature Summary and Revision History

Summary Data

Table 24: Summary Data

Applicable Product(s) or Functional Area	5G-UPF
Applicable Platform(s)	VPC-SI SMI
Feature Default Setting	Disabled – Configuration Required
Related Changes in this Release	Not Applicable
Related Documentation	*UCC 5G UPF Configuration and Administration Guide*

Revision History

Table 25: Revision History

Revision Details	Release
Support added for WiFi to LTE handover, and configuration to enable Converged Datapath feature at UPF.	2021.02.0

Revision Details	Release
First introduced.	2021.01.0

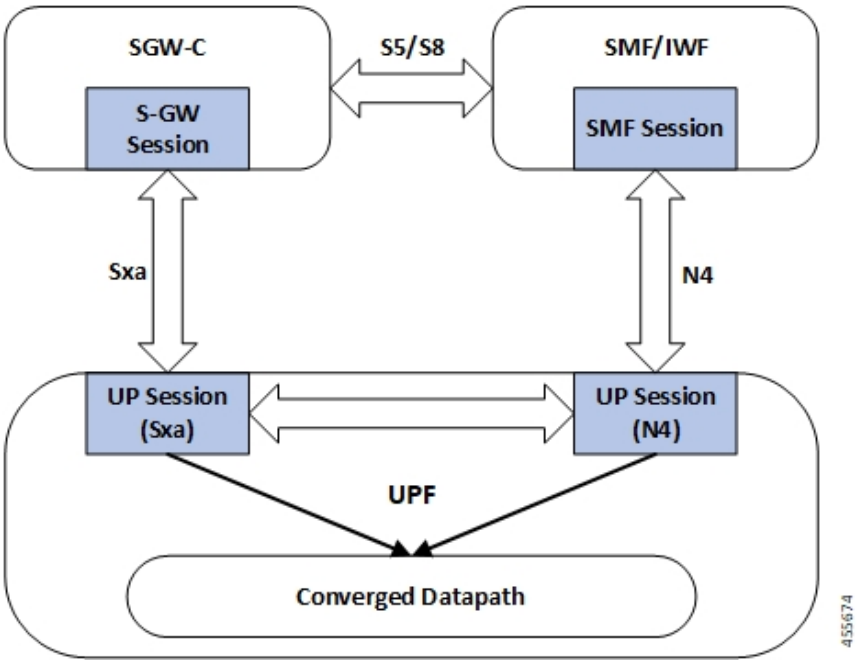
Feature Description

The Converged Datapath feature allows interconnection of the same UE's session at UPF instance with SGW-C/cnSGW and Session Management Function (SMF)/Inter-Working Function (IWF) to build converged/collapsed datapath and achieve higher throughput. With this feature:

- The UP/UPF selection logic is enhanced to aid same node selection on SGW-C/cnSGW and SMF.
- The SxDemux selects the same Session Manager (SessMgr) instance based on existing session of N4 or Sxa respectively.
- The Sxa session and N4 session correlation is done at SessMgr.
- The datapath is allowed to be collapsed in the forwarding plane.
- Extra hop in subscriber's datapath is eliminated, resulting in reduced latency and improved user experience.

Architecture

As part of this feature, there are two sessions on the same UPF instance established by SGW-C/cnSGW and SMF. Once they are established, the software logic determines the peer session so that the converged/collapsed datapath for packet processing is possible at the UPF node.



How it Works

This section describes how the feature works.

SxDemuxMgr

In distributed architecture of UP/UPF, sessions (Sxa or N4) run on different SessMgr instances. To support collapsing or converging the sessions to a single SessMgr, the SessMgr instance is selected by both sessions during establishment.

At SxDemux, when Sx Establishment Request (Sxa or N4) is received for selecting the SessMgr instance, it's parsed for finding the SessMgr instance from remote F-TEID, where corresponding sessions (N4 or Sxa respectively) are established. The F-TEID, that contains the Tunnel Identifier that is embedded with SessMgr instance, is extracted.

UPF also maintains IMSI entry at SxDemux. IMSI entry has information of SessMgr ID where the current session is hosted. When Converged Datapath feature is enabled, on receiving Sx Session Establishment Request, SxDemux first tries to find the SessMgr ID using F-TEID. If F-TEID is not present, IMSI lookup is done. If both F-TEID and IMSI is not present, then SxDemux doesn't select the same SessMgr ID for Sxa/N4 leg.

SessMgr

There are two sessions, Sxa and N4, that exist on the same SessMgr instance. To converge them, the following logic is used to identify the session:

- For Uplink Packet: Egress FARs F-TEID matches with Ingress PDRs F-TEID.
- For Downlink Packet: Ingress FAR's F-TEID should match with Egress PDR's F-TEID.

The F-TEID includes both Tunnel Identifier and the endpoint IP address. After the session is identified, the required information is used in datapath to build the converged datapath.

Datapath

After convergence of session occurs at SessMgr, the SessMgr removes the existing Bearer stream (3 tuple) from Fast Path that is installed for Sxa session. It's established only when flow-level stream (6 tuple), based on received packet, is analyzed.

The uplink packet is received by S-GW ingress PDR endpoint. The downlink packet is forwarded using S-GW Ingress FAR-based outer header.

Charging

UPF supports charging of SMF leg (N4 leg) and S-GW .

Note S-GW charging is not supported.

Call Flows

This section describes the call flows associated with Converged Datapath feature.

Initial Attach with SGW-C/cnSGW and SMF/IWF

The following illustration describes the initial attach call flow with collapsed UPF.

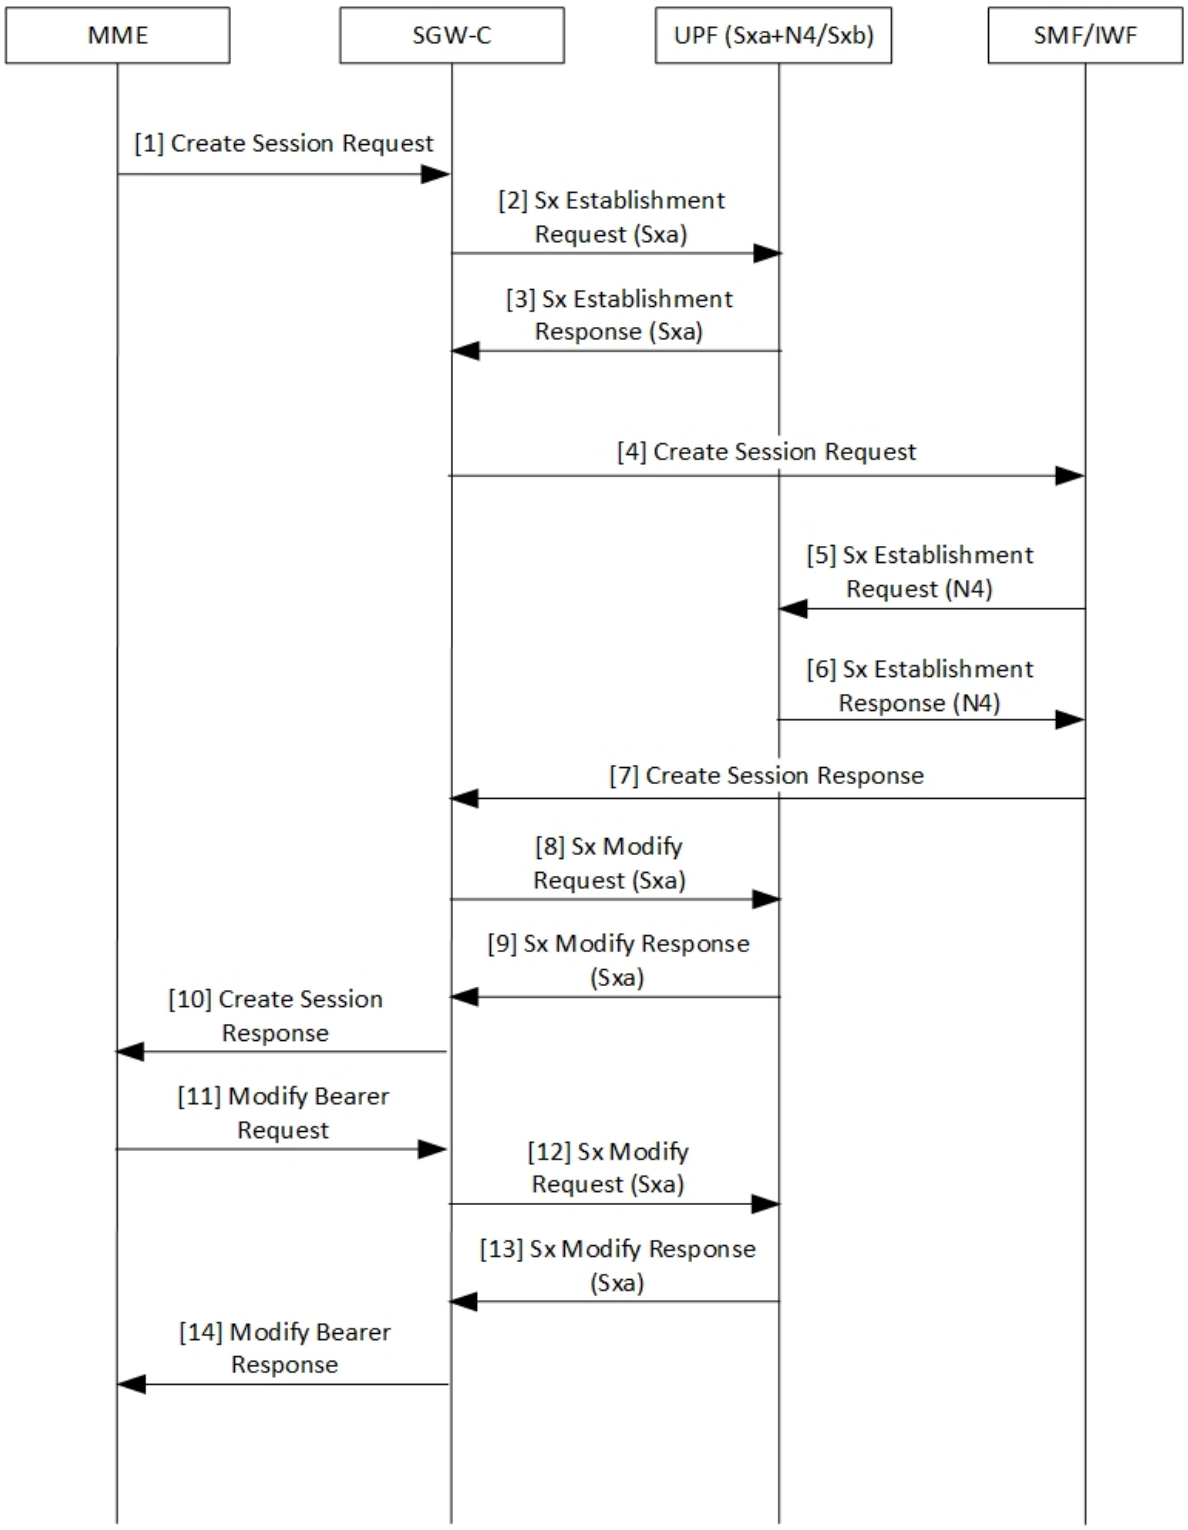

MME
SGW-C
UPF (Sxa+N4/Sxb)
SMF/IWF
[1] Create Session Request
[2] Sx Establishment Request (Sxa)
[3] Sx Establishment Response (Sxa)
[4] Create Session Request
[5] Sx Establishment Request (N4)
[6] Sx Establishment Response (N4)
[7] Create Session Response
[8] Sx Modify Request (Sxa)
[9] Sx Modify Response (Sxa)
[10] Create Session Response
[11] Modify Bearer Request
[12] Sx Modify Request (Sxa)
[13] Sx Modify Response (Sxa)
[14] Modify Bearer Response
455192

Table 26: Initial Attach with SGW-C/cnSGW and SMF/IWF Call Flow Description

Step	Description
1	Create Session Request (CSReq) is received by SGW-C/cnSGW and it selects the UPF.
2	The SGW-C/cnSGW sends Sx Establishment Request (Sxa) to the UPF. The UPF: • Allocates Sxa session • Allocates S-GW Ingress and Egress local F-TEID
3	The UPF sends Sx Establishment Response (Sxa) back to SGW-C/cnSGW.
4	The SGW-C/cnSGW sends CSReq to SMF/IWF. The SMF/IWF selects the same UPF that is selected by SGW-C/cnSGW.
5	The SMF/IWF sends Sx Establishment Request (N4) to the UPF.
6	The UPF sends Sx Establishment Response (N4) to the SMF/IWF.
7	The SMF/IWF sends Create Session Response to the SGW-C/cnSGW.
8	The SGW-C/cnSGW sends Sx Modify Request (Sxa) to the UPF. The UPF updates P-GW data F-TEID as part of Egress FAR. The UPF also interconnects Sxa and N4 session using internal logic and removes already-created Bearer Stream (3 tuple).
9	The UPF sends Sx Modify Response (Sxa) to the SGW-C/cnSGW.
10	The SGW-C/cnSGW sends Create Session Response to the MME.
11	The MME sends Modify Bearer Request to the SGW-C/cnSGW.
12	The SGW-C/cnSGW sends Sx Modify Request (Sxa) to the UPF. The UPF updates eNodeB F-TEID as part of Ingress FAR.
13	The UPF sends Sx Modify Response to the SGW-C/cnSGW.
14	The SGW-C/cnSGW sends Modify Bearer Response to the MME.

5G to 4G Handover with Collapsed UPF

The following illustration describes the 5G to 4G handover call flow with collapsed UPF.

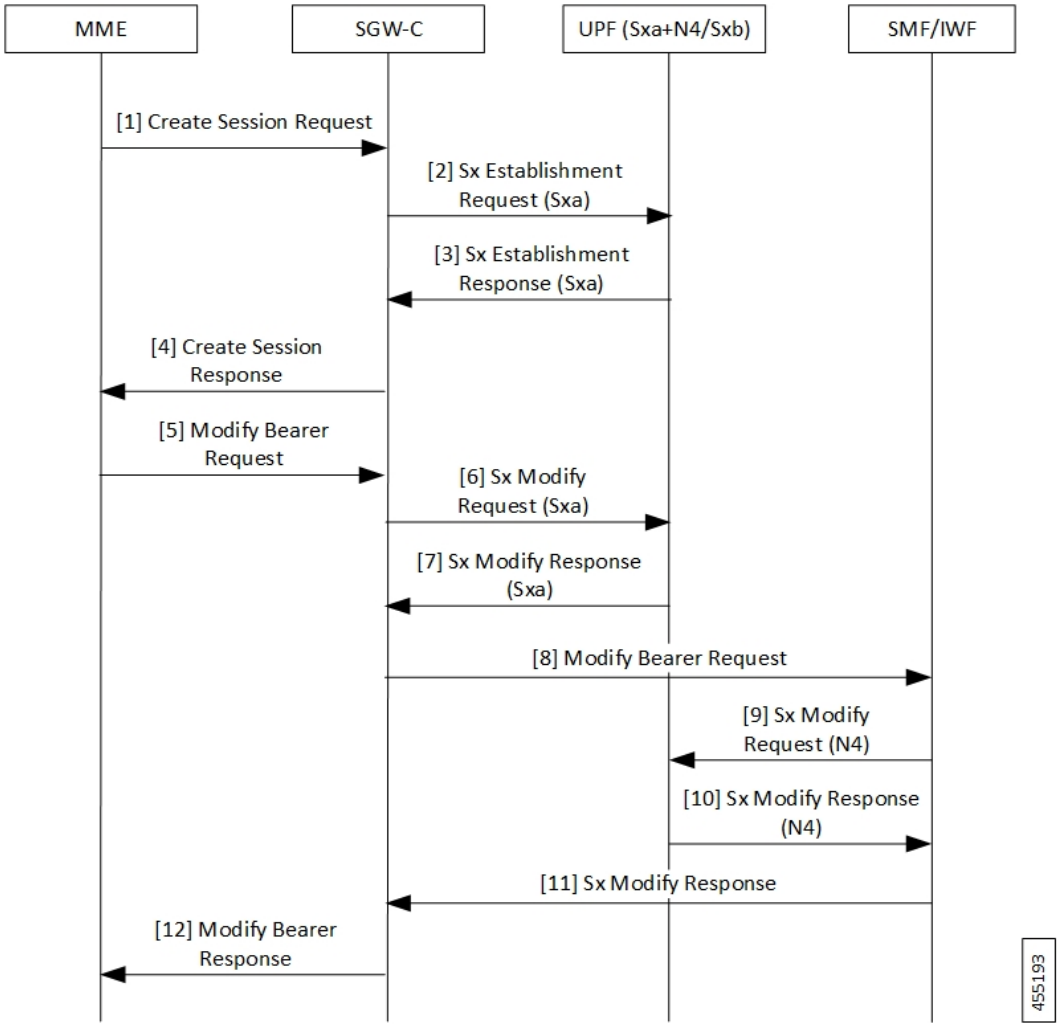


Table 27: 5G to 4G Handover with Collapsed UPF Call Flow Description

Step	Description
1	As part of UE initial attach, N4 session is already established with SMF and UPF. The MME sends Create Session Request (CSReq) to SGW-C/cnSGW. If it's a handover request, the SGW-C/cnSGW selects the same UPF that is selected by the SMF.
2	The SGW-C/cnSGW sends Sx Establishment Request (Sxa) to the UPF. At UPF: • SxDemux selects the same SessMgr instance extracted from the P-GW F-TEID that is received in FAR. Both Sxa and N4 session are on the same SessMgr. • Allocates Sxa session • Allocates S-GW Ingress and Egress local F-TEID • Interconnects Sxa and N4 session using internal logic and doesn't install Bearer Stream (3 tuple)

Step	Description
3	The UPF sends Sx Establishment Response (Sxa) back to SGW-C/cnSGW.
4	The SGW-C/cnSGW sends Create Session Response to the MME.
5	The MME sends Modify Bearer Request to the SGW-C/cnSGW.
6	The SGW-C/cnSGW sends Sx Modify Request (Sxa) to the UPF. The UPF updates eNodeB F-TEID for downlink data.
7	The UPF sends Sx Modify Response (Sxa) to the SGW-C/cnSGW.
8	The SGW-C/cnSGW sends Modify Bearer Request to the SMF/IWF.
9	The SMF/IWF sends Sx Modify Bearer Request (N4) to the UPF. The UPF: • Updates N4 session FAR towards S-GW with F-TEID • Updates TEP entries at VPP with new F-TEID
10	The UPF sends Sx Modify Response (N4) to the SMF/IWF.
11	The SMF/IWF sends Sx Modify Response to the SGW-C/cnSGW.
12	The SGW-C/cnSGW sends Modify Bearer Response to the MME.

Intra S-GW Handover with Collapsed UPF

The following illustration describes the intra-SGW handover call flow with collapsed UPF.

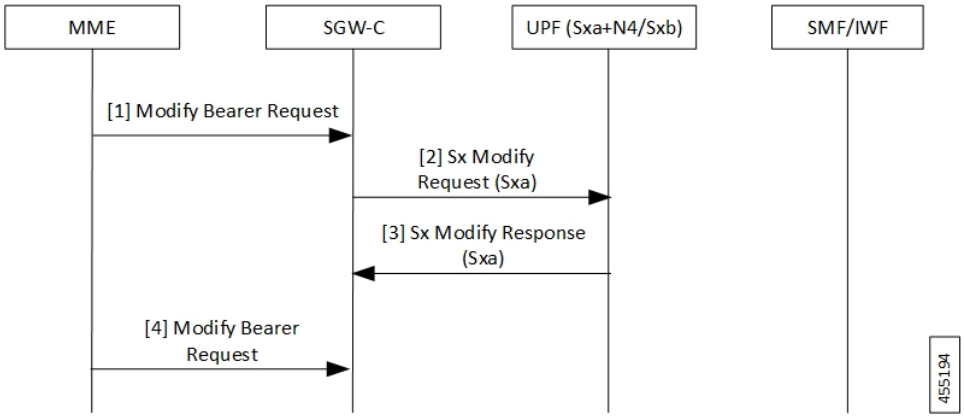


Table 28: Intra S-GW Handover with Collapsed UPF Call Flow Description

Step	Description
1	As part of UE initial attach, N4 and Sxa session is already established with SMF and UPF. At UPF, N4+Sxa session exists and are interconnected. The MME sends Modify Bearer Request to the SGW-C/cnSGW. The SGW-C/cnSGW updates eNodeB F-TEID in FAR.

Step	Description
2	The SGW-C/cnSGW sends Sx Modify Request (Sxa) to the UPF. The UPF: • Updates eNodeB F-TEID for downlink data • Updates TEP entries at VPP with new Remove TEID
3	The UPF sends Sx Modify Response (Sxa) to the SGW-C/cnSGW.
4	The SGW-C/cnSGW sends Modify Bearer Response to the MME.

Idle/Active DDN Handling with Collapsed UPF

The following illustration describes the Idle/Active DDN handling with collapsed UPF.

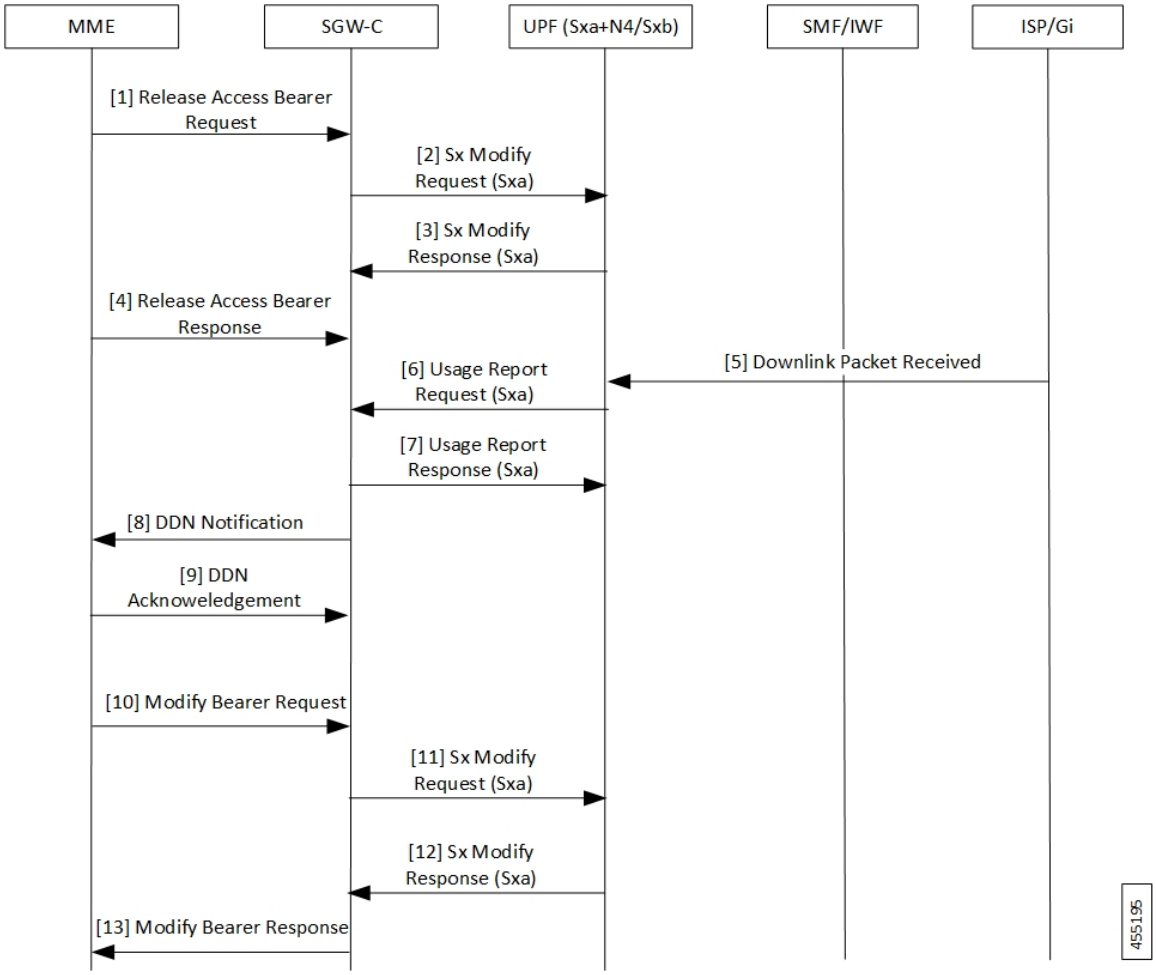


Table 29: Idle/Active DDN Handling with Collapsed UPF Call Flow Description

Step	Description
1	As part of UE initial attach, N4 and Sxa session is already established with SMF and UPF. At UPF, N4+Sxa session exists and are interconnected. At MME, the UE goes from Active to Ide state. The MME sends Release Access Bearer Request to the SGW-C/cnSGW. The SGW-C/cnSGW informs UPF to buffer packets.
2	The SGW-C/cnSGW sends Sx Modify Request (Sxa) to the UPF. The UPF updates FAR action towards eNodeB to buffer state. 6 tuple flows are onloaded for buffering.
3	The UPF sends Sx Modify Response (Sxa) back to SGW-C/cnSGW.
4	The MME sends Release Access Bearer Response to the SGW-C/cnSGW.
5	The ISP/Gi sends the received downlink packet to the UPF. The packet received by UPF at N4 session is passed to interconnect Sxa session for buffering.
6	The UPF sends Usage Report Request (Sxa) to the SGW-C/cnSGW.
7	The SGW-C/cnSGW sends Usage Report Response (Sxa) to the UPF.
8	The SGW-C/cnSGW sends DDN notification to the MME.
9	The MME sends DDN Acknowledgment back to the SGW-C/cnSGW.
10	At MME, the UE moves from Idle to Active. The MME sends Modify Bearer Request to the SGW-C/cnSGW. The SGW-C/cnSGW: • Updates eNodeB F-TEID in FAR • Updates FAR action to Forward
11	The SGW-C/cnSGW sends Sx Modify Request (Sxa) to the UPF. The UPF: • Updates eNodeB F-TEID for downlink data • Updates TEP entries at VPP with new Remove TEID • Releases buffered packets by finding respective 6 tuple streams
12	The UPF sends Sx Modify Response (Sxa) to the SGW-C/cnSGW.
13	The SGW-C/cnSGW sends Modify Bearer Response to the MME.

IDFT Handling during S1 Handover

The following illustration describes the IDFT handling during S1 handover with collapsed UPF.

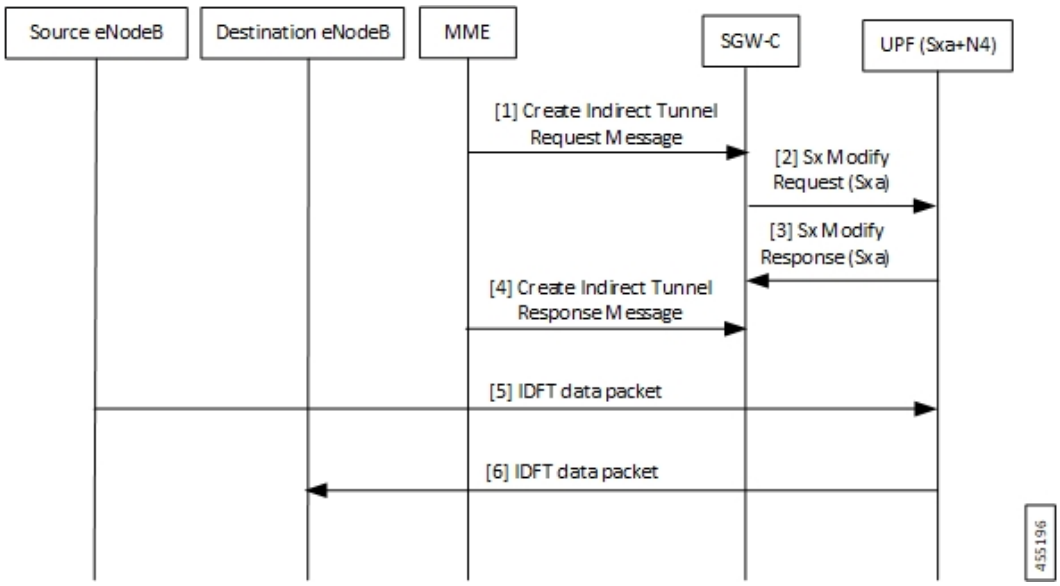


Table 30: IDFT Handling during S1 Handover Call Flow Description

Step	Description
1	As part of initial attach, N4 and Sxa session is already established with SMF and UPF. S1 handoff is triggered. The MME sends Indirect Tunnel Request Message to the SGW-C/cnSGW.
2	The SGW-C/cnSGW sends Sx Modify Request (Sxa) to the UPF. At UPF: • IDFT PDR is detected at SessMgr • New F-TEID is allocated for IDFT tunnel • Converged datapath is not required and traffic goes through Slowpath
3	The UPF sends Sx Modify Response (Sxa) to the SGW-C/cnSGW.
4	The SGW-C/cnSGW sends Create Indirect Tunnel Response Message to the MME.
5	The Source eNodeB sends IDFT data packet to the UPF. If there's no matching 3 tuple stream at UPF, the packet is forwarded to SessMgr.
6	The UPF sends IDFT data packet to the destination eNodeB.

S-GW Relocation with Same SGW-U

The following illustration describes the S-GW relocation with destination S-GW selecting the same UPF.

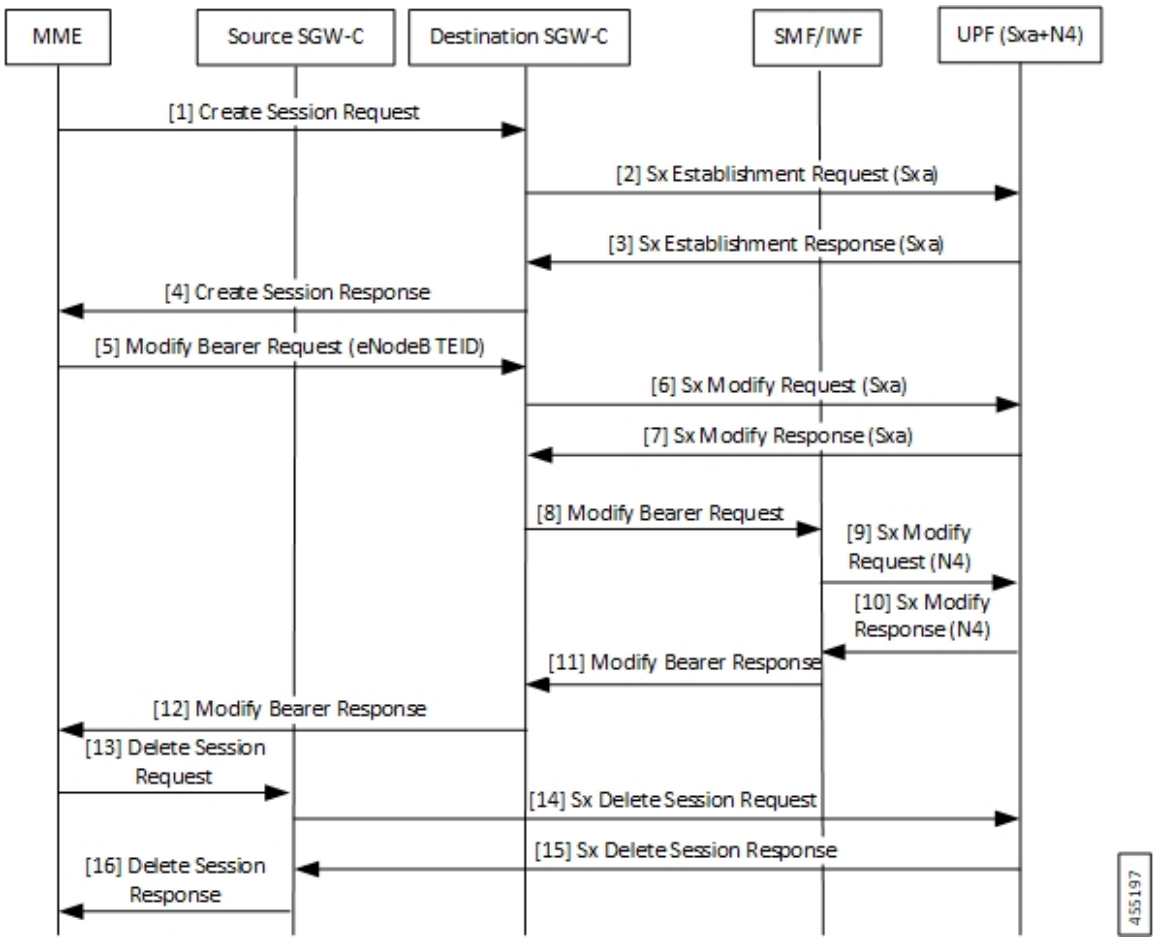


Table 31: S-GW Relocation with Same SGW-U Call Flow Description

Step	Description
1	As part of initial attach, N4 and Sxa session is already established with Source SGW-C/cnSGW, SMF/IWF, and UPF. The MME sends Create Session Request to the destination SGW-C/cnSGW.
2	The destination SGW-C/cnSGW sends Sx Establishment Request (Sxa) to the UPF. The UPF links new Sxa session with the N4 session for uplink packets (Slowpath).
3	The UPF sends Sx Establishment Response (Sxa) to the destination SGW-C/cnSGW.
4	The destination SGW-C/cnSGW sends Create Session Response to the MME.
5	The MME sends Modify Bearer Request (eNodeB F-TEID) to the destination SGW-C/cnSGW.
6	The destination SGW-C/cnSGW sends Sx Modify Request (Sxa) to the UPF.
7	The UPF sends Sx Modify Response (Sxa) to the destination SGW-C/cnSGW.
8	The destination SGW-C/cnSGW sends Modify Bearer Request to the SMF/IWF.

Step	Description
9	The SMF/IWF sends Sx Modify Request (N4) to the UPF. The UPF switches downlink tunnel and links the N4 session with the new Sxa session.
10	The UPF sends Sx Modify Response (N4) to the SMF/IWF.
11	The SMF/IWF sends Modify Bearer Response to the destination SGW-C/cnSGW.
12	The destination SGW-C/cnSGW sends Modify Bearer Response to the MME.
13	The MME sends Delete Session Request to the source SGW-C/cnSGW.
14	The source SGW-C/cnSGW sends Sx Delete Session Request to the UPF.
15	The UPF sends Sx Delete Session Response to the source SGW-C/cnSGW.
16	The source SGW-C/cnSGW sends Delete Session Request to the MME.

WiFi to LTE Handover

The following illustration describes the WiFi to LTE handover call flow.

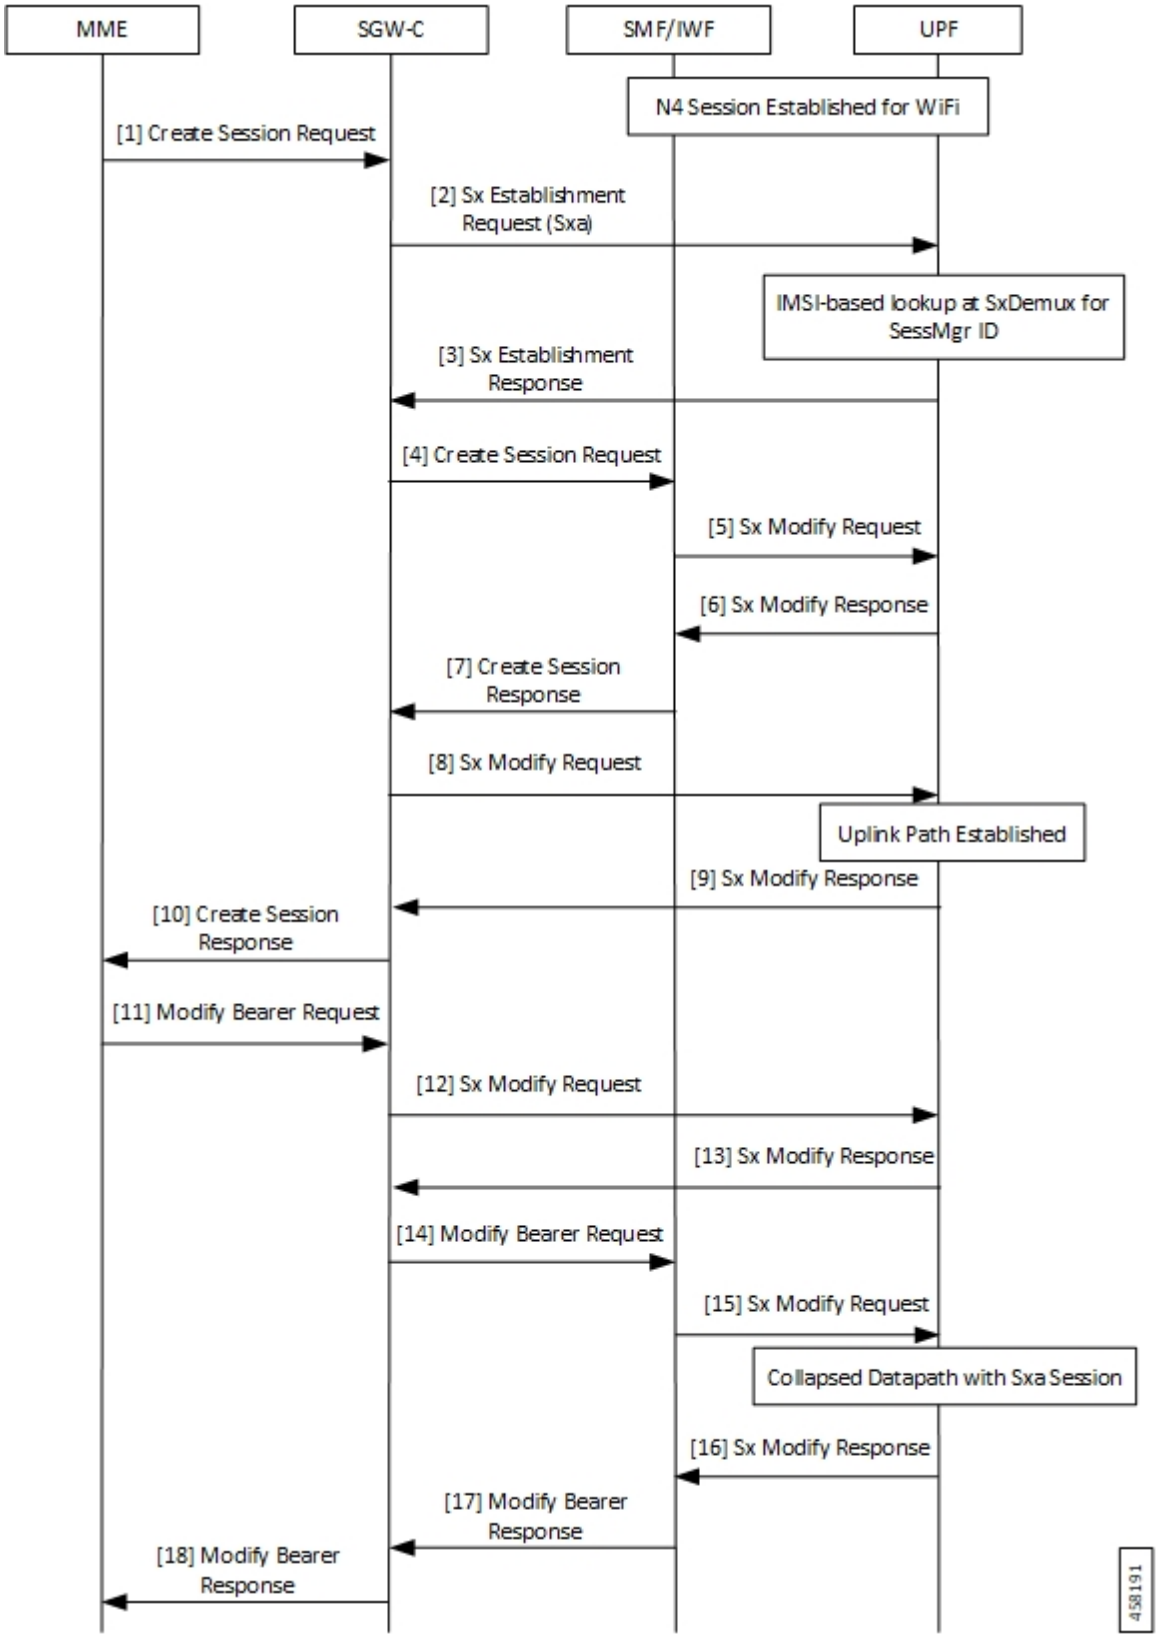

MME
SGW-C
SMF/IWF
UPF
N4 Session Established for WiFi
[1] Create Session Request
[2] Sx Establishment Request (Sxa)
IMSI-based lookup at SxDemux for SessMgr ID
[3] Sx Establishment Response
[4] Create Session Request
[5] Sx Modify Request
[6] Sx Modify Response
[7] Create Session Response
[8] Sx Modify Request
Uplink Path Established
[9] Sx Modify Response
[10] Create Session Response
[11] Modify Bearer Request
[12] Sx Modify Request
[13] Sx Modify Response
[14] Modify Bearer Request
[15] Sx Modify Request
Collapsed Datapath with Sxa Session
[16] Sx Modify Response
[17] Modify Bearer Response
[18] Modify Bearer Response
458191

Table 32: WiFi to LTE Handover Call Flow Description

Step	Description
1	N4 Session is established for WiFi between SMF/IWF and UPF. The MME sends Create Session Request (relocating to 4G) to the SGW-C/cnSGW.
2	The SGW-C/cnSGW sends Sx Session Establishment Request (No F-TEID) to the UPF.
3	IMSI-based lookup is done at SxDemux for SessMgr ID. The UPF sends Sx Establishment Response to the SGW-C/cnSGW.
4	The SGW-C/cnSGW sends Create Session Request to the SMF/IWF.
5	The SMF/IWF sends Sx Modify Request with new Uplink RB PDR to the UPF.
6	The UPF sends Sx Modify Response to the SMF/IWF with new TEID for Uplink data.
7	The SMF/IWF sends Create Session Response to the SGW-C/cnSGW.
8	The SGW-C/cnSGW sends Sx Modify Request to the UPF with updated Uplink FAR.
9	Uplink path is established. The UPF sends Sx Modify Response to the SGW-C/cnSGW.
10	The SGW-C/cnSGW sends Create Session Response to the MME.
11	The MME sends Modify Bearer Request to the SGW-C/cnSGW.
12	The SGW-C/cnSGW sends Sx Modify Request to the UPF with eNodeB TEIDs.
13	The UPF sends Sx Modify Response back to the SGW-C/cnSGW.
14	The SGW-C/cnSGW sends Modify Bearer Request to the SMF/IWF.
15	The SMF/IWF sends tunnel switch Sx Modify Request to the UPF.
16	Collapsed datapath with Sxa session is established at UPF. The UPF sends Sx Modify Response to the SMF/IWF.
17	The SMF/IWF sends Modify Bearer Response to the SGW-C/cnSGW.
18	The SGW-C/cnSGW sends Modify Bearer Response to the MME.

Limitations

The following are the known limitations of the feature:

- If Sxa leg is of one user-plane-service and N4 leg is of another user-plane-service, then datapath won't be collapsed.
- If Sxa leg is under one context and N4 leg is in another context, then datapath can't be collapsed.
- The S-GW part of the **show subscribers user-plane-only full all** CLI output doesn't display ToS-marked Uplink and Downlink packets.

- For the S-GW part of the **show user-plane-service statistics rat all** CLI output, the session statistics for Unknown is incremented, however, the data statistics aren't incremented under RAT-type Unknown.
- Lawful Intercept for S-GW isn't supported.
- S-GW charging isn't supported.
- S-GW bearer inactivity timeout isn't honored, as it's determined by S-GW URR for which processing isn't done for collapsed datapath.
- If S-GW leg of call is locally purged for converged session, then the UPF continues to send data toward eNodeB.

Charging for Converged Calls

Table 33: Feature History

Feature Name	Release Information	Description
Charging Support for Converged Calls	2024.02.0	UPF supports S-GW Charging Data Record (CDR) volume reporting for converged calls. This feature allows you to use these CDRs for reporting and charging during inbound LBO roaming scenario. This feature introduces a new CLI command **converged-sxa-usage-reporting** in the ACS Service configuration, to generate the usage report with volume count. **Default Setting**: Disabled – Configuration Required to Enable

Feature Description

For a converged call, SGW-U sends the usage report to SGW-C based on which the SGW-C generates the CDR. You can enable this feature through the **converged-sxa-usage-reporting** CLI command. Once enabled, SGW-U sends a usage report with volume accounting to SGW-C, based on which the SGW-C generates the CDR.

How it Works

For the SGW-U to update a volume report, the URR linking happens between SGW-U bearer URR and PGW-U, UPF bearer, or SDF level URRs. While the reconciliation of charging on UPF, the reconciliation of charging on Sxa also happens simultaneously. Depending on the threshold triggers, the SGW-U generates the usage report. Examples of threshold triggers are time threshold, volume threshold, and Query URR. Then, SGW-U sends the usage report to SGW-C based on which the SGW-C generates CDR.

To update the volume count for SGW-U bearer URRs along with PGW-U or UPF URRs for a converged call, the following process happens:

- The reference pointer of SGW-U default bearer URR is stored at PGW-U or UPF default bearer URR.
- The reference pointer of SGW-U dedicated bearer URRs is stored at SDF level URRs of the corresponding dedicated bearer of PGW-U or UPF.

While reconciling the LC count, the URRs of PGW-U or UPF and SGW-U bearer are traversed. Then, the volume counts are updated for both the call ID URRs.



Note

- If a call is established as a converged call, then the reference pointer of SGW-U bearer URR is stored at PGW-U, UPF bearer, or SDF level URR while call convergence.
- In the case of S-GW relocation from a nonconverged call to a converged call, the URR linking happens. The URR linking implies that the reference pointer of the SGW-U bearer URR is stored at PGW-U, UPF bearer, or SDF level URR.
- In the case of S-GW relocation from a converged call to a nonconverged call, the URR delinking happens. The URR delinking implies that the reference pointer of SGW-U bearer URR is removed from PGW-U, UPF bearer, or SDF level URR.
- In the case of recovery, recreation of the link happens, which implies that the reference pointer of SGW-U URR is stored at PGW-U or UPF URR during recovery to maintain consistent behavior.

Limitations

The charging for converged calls feature has the following limitations:

- Volume reporting in S-GW URRs is based on the corresponding PGW-U or UPFs URR usage count. For example, the uplink packets that are dropped due to P-GW processing aren’t accounted in the Sxa call.
- RA or RS packet charging doesn't happen at UPF. As the Sxa usage report depends on the UPF, the RA or RS packets aren't accounted for charging on the Sxa call.
- Multiple bearers using the same URR ID lead to inconsistency in the usage report that is generated as Sxa.
- For converged calls, if packets aren’t charged at N4 leg, then these packets aren't charged at Sxa as well.
- Charging for predefined rules on a dedicated bearer isn't supported for converged calls.
- When the call is getting converged, the feature configuration, either enabled or disabled, remains the same until the duration of the existing call. The change in configuration is applicable from the new call onwards. If the feature is enabled, the charging continues for existing and new URRs. However, if the feature is disabled, the charging doesn't happen for the existing and new URRs.

Configuring eGCDR

The eGCDR configuration is required on PGW-U or UPF to create static default bearer URR at PGW-U or UPF.

The following is a sample configuration:

```
config
   active-charging service ACS
       rulebase starent
           billing-records egcdr
   end
```

Configuring GTPP Group

The GTPP group configuration is required on PGW-U or UPF to create static default bearer URR at PGW-U or UPF.

The following is a sample configuration:

```
config
   context ISP
      apn intershat
          gtpp group my_grp
   end
```

Generating Usage Report with Volume Count for Converged Sxa Calls

To generate the usage report with volume count for Sxa call, use the following sample configuration:

```
config
   active-charging service service_name
      converged-sxa-usage-reporting
         end
```

NOTES:

- **converged-sxa-usage-reporting**: Enable generation of usage reports for converged SxA calls. To disable this functionality, add the **no converged-sxa-usage-reporting** syntax. The default value is **no converged-sxa-usage-reporting**.

Configuration Example

The following is an example configuration of generation of usage report for the converged Sxa call:

```
config
  active-charging service ACS
     converged-sxa-usage-reporting
  end
```

Configuring Converged Datapath

This section describes the CLI commands available in support of the feature.

Enabling Converged Datapath at UPF

Use the following configuration to enable Converged Datapath at UPF.

```
configure
   user-plane converged-mode
   end
```

NOTES:

- **user-plane**: Specifies the UPF related to the configuration.
- **converged-mode**: Specifies the collapse datapath of Sxa and N4 calls.
- By default, the CLI is disabled.
- It is recommended to add the CLI in boot configuration.

Configuring Remote Peers for Sxa and N4

Use the following CLI commands to configure remote peers for Sxa and N4 interface.

```
configure
  control-plane-group group_name
    peer-node-id ipv4-address ipv4_address interface n4
    peer-node-id ipv4-address ipv4_address
    end
```

Configuring User Plane Service for Sxa and N4

Use the following CLI commands to configure UP Service for Sxa and N4 interface.

```
configure
  user-plane-service service_name
    associate gtpu-service gtpu_service upf-ingress
    associate gtpu-service gtpu_service sgw-ingress
    associate gtpu-service gtpu_service sgw-egress
    associate gtpu-service gtpu_service cp-tunnel
    associate sx-service sx_service
    associate control-plane-group group_service
    end
```

Monitoring and Troubleshooting

Show Commands and/or Outputs

This section provides information about the show CLI commands and/or outputs available in support of the Converged Datapath feature.

show subscribers user-plane-only all

The output of this CLI command is enhanced to display the following fields to indicate if the session is converged or non-converged:

- Converged
- Non Converged

show subscribers user-plane-only full all

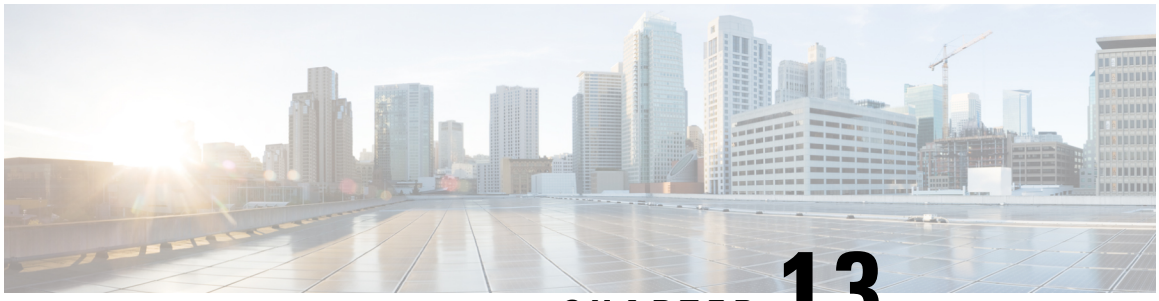
The output of this CLI command is enhanced to display the following fields:

- Converged Session
- Converged Peer Callid

show user-plane-service statistics all

The output of this CLI command is enhanced to display the following fields:

- Converged Data Session PDNs:
  - Active
  - Setup
  - Released



CHAPTER 13

Deep Packet Inspection and Inline Services

- Feature Summary and Revision History, on page 121
- Feature Description, on page 122
- How it Works, on page 124
- Supported Inline Services, on page 127
- Configuring the Static and Pre-Defined Rules, on page 183
- Configuring ACS Ruledef for L7 Protocols for DPI, on page 184
- Charging Action Configuration for L7 Protocols for DPI, on page 186

Feature Summary and Revision History

Summary Data

Applicable Product(s) or Functional Area	5G-UPF
Applicable Platform(s)	VPC-SI SMI
Feature Default Setting	Disabled – Configuration Required
Related Changes in this Release	Not Applicable
Related Documentation	*UCC 5G UPF Configuration and Administration Guide*

Revision History

Revision Details	Release
• NAT Support on the N4 interface has been added. • Subscriber Firewall support in a TCP Idle Timeout has been added. • Support added for capturing data stall issues using 2 new attributes.	2023.04.0

Revision Details	Release
Support has been added for the following functionality: • IP Readdressing • RTP Dynamic Flow Detection • Rule-matching for Bearer-specific Filters • QUIC IETF implementation	2021.02.0
New L7 protocols have been introduced as part of Deep Packet Inspection (DPI).	2021.01.0
The following EDR attributes have been added for TCP: • SYN and SYN-ACK packet • SYN-ACK and ACK packet	2021.01.0
New DNS attributes have been introduced in EDRs.	2021.01.0
First introduced.	2020.02.0

Feature Description

One of the key product capability of Cisco 5G-UPF is integrated Deep Packet Inspection (DPI) based services. DPI is the examination of layer 7 (L7), which contains Uniform Resource Identifier (URI) information. In some cases, layer 3 (L3) and layer 4 (L4) analyzers that identify a trigger condition are insufficient for billing purposes, so layer 7 (L7) examination is used.

DPI performs packet inspection beyond L4 inspection and is typically deployed for detection of URI information at L7 (for example, DNS, HTTP, HTTPS, RTP, and RTSP URLs).

DNS Server Readdressing

Whenever you use an unauthorized DNS server, you can modify the request to readdress the DNS IPs to use authorized servers. A ruledef determines if a packet belongs to a DNS query. It also determines if the DNS query belongs to a set of authorized DNS servers. If the DNS query does not belong to the authorized DNS servers, the flow action picks up DNS servers from the readdress server list.

You can configure the **readdress-server-list** command under active-charging service. When the flow matches a **ruledef**, you can configure the flow action to use the servers from **readdress-server-list**.

To configure the readdress server list under active-charging service, use the following configuration:

```
configure
  active-charging service service_name
    readdress-server-list name_of_list
      server ipv4_address [ port ]
      server ipv6_address [ port ]
```



Note You can configure a maximum number of 10 servers in a readdress server list and a maximum of 10 readdress server lists under active-charging service. Both IPv4 and IPv6 addresses can be configured in the same **readdress-server-list**.

Select the **readdress-server-list** from the list using one of the following methods:

- Round-robin—Server selection occurs in a round-robin manner for every new flow. Inactive servers in the list are not considered during the selection.

  This method is the default selection.

- Hierarchy—The servers that are tagged in this approach are primary, secondary, tertiary, and so on, depending on the order they are defined in the **readdress-server-list**. All flows are readdressed to the primary server as long as it is available. If the primary server goes down, then flows are readdressed to the secondary server and the same logic recurs. Once the primary server is active, then flows switch back to the primary server for readdressing.

To configure the DNS readdress server list, use the following CLI configuration under active-charging service.

```
configure
  active-charging service service_name
    readdress-server-list name_of_list
      server ipv4_address [ port ]
      server ipv6_address [ port ]
      consecutive-failures integer_value
      response-timeout integer_value
      reactivation-time integer_value

      charging-action action_name
        flow action readdress server-list name_of_list
        exit
```

NOTES:

- **consecutive-failures**—Specify an integer ranging from 1–10. The default value is *5*.
- **response-timeout**—Specify an integer ranging from 1–10000 milliseconds. The default value is *1000*.
- **reactivation-time**—Specify an integer ranging from 1–1800 seconds. The default value is *300*.

Readdress Server States

This section describes the readdress server states:

- Active state—Once configured, all servers are marked as Active.
- Inactive state—If no response is received from the readdressed server, then the server is marked as Inactive.
- Active-Pending state—Once the server is in Active-Pending state, it is available to accept the requests for readdressing. In this state, if a request is readdressed to this server and response is returned from it, then the server state is changed to Active. Otherwise, it is moved back to Inactive state.

How it Works

This section describes the following functionality of DPI:

- DSCP Marking of downlink and uplink packets.
- Traffic Readdressing or Redirecting.

DSCP Marking for Downlink and Uplink Packets

Table 34: Feature History

Feature Name	Release Information	Feature Description
EDR Attributes for DSCP Mapping	2024.02.0	UPF supports two new EDR attributes **sn-dscp-uplink** and **sn-dscp-downlink** to mark the user plane traffic using the dscp value received from SMF or DSCP(derived from TOS) configured under the charging-action. These attributes allow the users to understand the QoS flow information of the uplink and downlink traffic in their network. **Default Setting**: Disabled — Configuration Required to Enable

Transport-level marking is the process of marking traffic at UPF with a Differentiated Services Code Point (DSCP) value. The transport-level marking, executed on per-QoS flow, is based on the mapping from 5QI and optional Allocation and Retention Policy (ARP) configuration from SMF.

SMF controls the transport-level marking by providing the DSCP in the ToS (IPv4) or Traffic Class (IPv6) within the Transport Level Marking IE in the FAR, that is associated to the PDR matching the traffic to be marked. UPF performs the transport level marking for detected traffic and sends the marked packet to the peer entity. The SMF can change the transport-level marking by changing the Transport Level Marking IE in the related FAR.

UPF supports the following IEs:

- Inner Packet Marking IE—Marks the tunnel packets for inner packet marking.
- Transport Level Marking Options IE—Copies DSCP of the inner packet to the outer IP header.

EDR Attributes

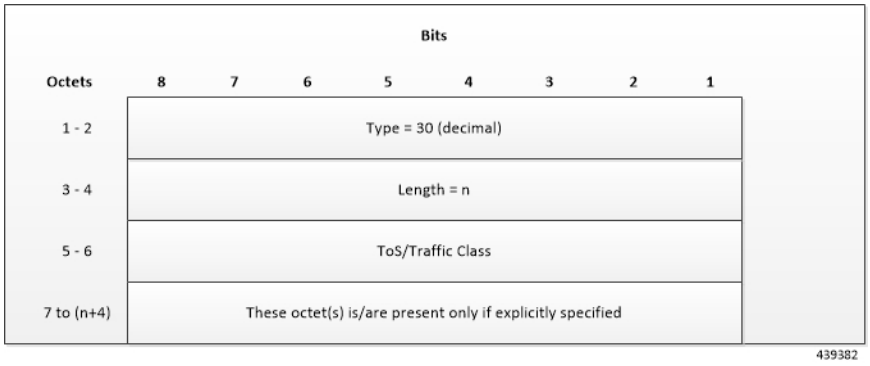
UPF records the DSCP values marked on the last outgoing traffic using the EDR attributes **sn-dscp-uplink** and **sn-dscp-downlink**.

UPF derives the DSCP values from the **ip tos** configuration under ACS charging-action or Transport Level Marking AVP received in FAR from SMF. The DSCP value from the IP ToS configuration takes higher precedence over that from the Transport Level Marking AVP.

For more information, see the *Event Data Records* section.

Transport Level Marking IE

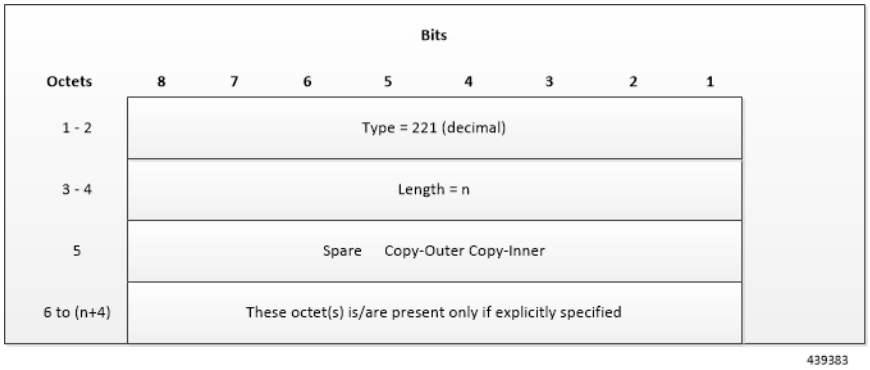
The "Transport Level Marking" IE type is encoded as shown in the following figure. It indicates the DSCP value for the downlink transport-level marking.



The encoding for Type-of-Service (ToS) or Traffic Class takes place in the form of two octets as an OctetString. The first octet contains the DSCP value in the IPv4 Type-of-Service or the IPv6 Traffic-Class field and the second octet contains the ToS/Traffic Class mask field, which is set to *0xFC*.

Transport Level Marking Options IE

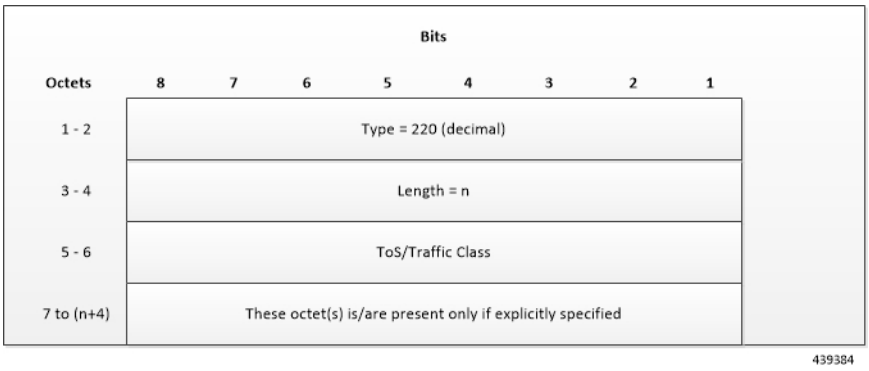
The "Transport Level Marking Options" IE type is encoded as shown in the following figure. The DSCP for downlink transport-level marking is copied from the inner packet.



The Copy-Inner and Copy-Outer flags are present in bit-0 and bit-1 of octet 5. Copy-Outer flag is not used for downlink packets because there is no outer header present in packets coming from ISP. If a Copy-Inner flag is present, then the UPF uses DSCP value from the inner packet to mark the transport-level IP header.

Inner Packet Marking IE

The "Inner Packet Marking" IE type is encoded as shown in the following figure. It indicates the DSCP value for the downlink inner packet marking.



The encoding for ToS/Traffic Class takes place in the form of two octets as an OctetString. The first octet contains the DSCP value in the IPv4 ToS or the IPv6 Traffic Class field and the second octet contains the ToS/Traffic Class mask field, which is set to *0xFC*.

NOTES:

- The original ECN bits in the IP header of User Plane packets do not change after applying transport-level marking or inner packet marking.
- If "Transport Level Marking" IE, "Inner Packet Marking" IE, or both the IEs are associated with uplink FAR, then the following rule applies for uplink packet marking:
  - If "Transport Level Marking" or "Inner Packet Marking" IE is present, its DSCP value is used.
  - If both "Transport Level Marking" and "Inner Packet Marking IE" are present, then the value from "Transport Level Marking" IE is used for uplink packet marking.

Traffic Readdressing or Redirecting

Traffic Redirection is the process of redirecting uplink application traffic to a redirect destination. For example, redirect some HTTP flows to service provisioning page. The redirect destination is provided by the PCF or it is preconfigured in the SMF or in the UPF.

The traffic redirection enforcement is applicable for the SMF or in the UPF if the traffic that the UPF supports subjects to traffic redirection. The UPF reports to the SMF whether it supports traffic redirection enforcement in the UPF through the "UP-Function Features" IE.

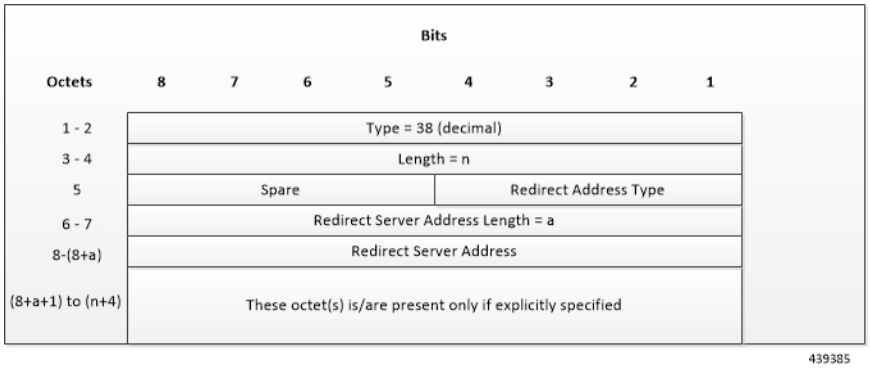
To enforce the traffic redirection in the UPF, the SMF takes the following actions:

- Creates the necessary PDRs, if it does not exist, to represent the traffic redirection.
- Creates a FAR with:
  - The "Redirect Information" IE, that includes the redirect destination, if the traffic needs redirection toward a redirect destination that is provided by the SMF. The redirect destination from the SMF prevails over a redirect destination that is preconfigured in the UPF.
    - For HTTP traffic redirection, the Redirection Address Type is set to "URL" and the SMF sets the "Destination Interface" IE in the FAR to "Access" (to forward the HTTP Response message with a status-code indicating redirect). For other types of traffic redirection, the "Destination Interface" IE in the FAR is set to "Core".

- Associates the FAR to the above PDRs of the PFCP session.

Redirect Information IE

Redirect information is encoded as follows:



"Redirect Address Type" indicates the type of the redirect address:

Redirect Address Type	Value (Decimal)
IPv4 address	0
IPv6 address	1
URL	2
SIP URI	3
Spare, for future use.	4–15

The "Redirect Server Address Length" indicates the length of the "Redirect Server Address". The "Redirect Server Address" encoding is in UTF8String format and contains the address of the redirect server (for example, HTTP redirect server, SIP server) with which the end user connects.

Important In this release, only Redirect Address Type URL is supported for dynamic rule when FAR is associated with URR where quota expires.

Supported Inline Services

Application Detection and Control

The ADC in-line service is used to detect Peer-to-Peer protocols by analyzing traffic. Other popular applications that generate the bulk of Internet traffic like Social Networking and Gaming applications can be detected.

The in-line service known as ADC is also referred as "P2P". Peer to Peer (P2P) is a term used in two slightly different contexts. At a functional level, it means protocols that interact in a peering manner, in contrast to client/server manner. There is no clear differentiation between the function of one node or another. Any node can function as a client, a server, or both—a protocol may not clearly differentiate between the two. For

example, peering exchanges may simultaneously include client and server functionality, sending and receiving information.

Note The ADC support is a licensed feature. Contact your Cisco Account or Support representative for information on how to obtain a license.

QUIC IETF Implementation

In the current framework, Deep Packet Inspection (DPI) is done for every packet in a flow when it reaches the plugin. The DPI is done by analyzing the packets and extracting deterministic patterns. The DPI is done in-order to detect the application and to classify its subtype. Plugin excludes the flow after the DPI. The flow is offloaded after the detection. As part of QUIC IETF, the initial QUIC handshake packets (Client/Server Hello) are encrypted over the network. Hence, there are no deterministic patterns available for detection of the application. Support is added in p2p plugin to decrypt and obtain the SNI (Server Name Indication) for detection.

Configuring QUIC IETF

Use the following configuration to enable or disable the QUIC IETF decryption.

```
configure
  active-charging service acs_service_name
    p2p-detection debug-param protocol-param p2p_quic_ietf_decrypt 1
    end
```

Note

- By default, the CLI is disabled and there's minimal impact on the performance due to TLS decryption.
- Runtime change of configuration doesn't impact the existing flow. Change is applicable only for new flow.

Statistics

show user-plane-service statistics analyzer name p2p

Use this show CLI command to determine the packets that are analyzed for QUIC and application.

SCONE support

Table 35: Feature history

Feature name	Release	Feature description
SCONE protocol for UPF	2026.03.0	This release introduces support for the SCONE protocol in the User Plane Function (UPF), enabling network-aware throughput optimization for QUIC-based applications. SCONE allows the UPF to provide explicit throughput advice to SCONE-capable endpoints by updating a rate signal in QUIC SCONE packets, helping applications adapt to network-enforced bandwidth policies without excessive congestion probing.

Modern applications that use QUIC often cannot determine the throughput constraints enforced by the network. As a result, congestion control algorithms continue probing for additional bandwidth, which can lead to unnecessary packet loss, increased latency, and inefficient bandwidth utilization. SCONE (Standard Communication with Network Elements) enables the UPF to communicate throughput advice to SCONE-capable QUIC endpoints by modifying a 7-bit rate signal field in QUIC transport parameters. This allows applications to optimize their transmission rate based on network policies rather than relying solely on congestion detection, making network-enforced throughput policies visible and reducing unnecessary congestion probing.

SCONE provides the following benefits:

- Makes network-enforced throughput policies visible to QUIC endpoints.
- Reduces unnecessary congestion probing performed by congestion control algorithms.
- Enables applications to optimize transmission rates based on network guidance instead of packet loss or delay.
- Improves bandwidth utilization by communicating explicit throughput advice to SCONE-capable endpoints.
- Enables network-aware traffic optimization without requiring application-specific modifications.

SCONE Operational Flow

SCONE processing consists of the following stages:

- **Flow Detection:** ADC-P2P identifies SCONE-capable QUIC flows by inspecting the QUIC Transport Parameters for the `scone_supported` transport parameter (ID `0x219e`) and verifying the SCONE indication signature (`0xC8 0x13`) in the QUIC Initial packet. If both identifiers are present, the flow is marked as SCONE-capable for its lifetime. Detection is performed only during flow establishment and is not applied to existing flows when SCONE is enabled mid-session.

- **Rule Matching:** The `scone-detected` ruledef keyword allows operators to distinguish SCONE-capable flows from other traffic and apply dedicated charging actions based on SCONE detection status.
- **Rate Signal Calculation:** The UPF determines the throughput advice using either an explicitly configured `ta-value` or the most restrictive value between the flow Maximum Bit Rate (MBR) and the APN-AMBR. For dynamic rules, the MBR is obtained from the QER associated with the matching PFCP PDR. The calculated bitrate is encoded into a 7-bit SCONE Rate Signal.
- **Rate Signal Modification:** For every downlink SCONE packet, the VPP validates the SCONE packet format, extracts the existing Rate Signal, and compares it with the value calculated by the UPF. The packet is modified only when the calculated value is lower than the Rate Signal already present in the packet. After modification, the UDP checksum is recalculated before forwarding the packet.
- **EDR Reporting:** When configured, the UPF reports SCONE-specific information in the Charging EDR, including the SCONE detection status (`scone-enabled`) and the last Rate Signal value (`scone-rate-signal`) applied to the flow.

Configure SCONE support

SCONE enables the UPF to detect SCONE-capable QUIC flows and communicate throughput advice to supported QUIC endpoints based on configured network policies.

Before you begin

- Ensure that the UPF is running a software release that supports the SCONE feature.
- Enable QUIC detection in the ADC service.
- Enable P2P dynamic flow detection on the rulebase before configuring SCONE.
- Ensure that the required rulebase, ruledefs, and charging actions are created.

Follow these steps to configure SCONE support:

Procedure

Step 1 Enable SCONE detection in the ADC service:

```
p2p-detection protocol all quic
p2p-detection debug-param protocol-param p2p_quic_ietf_decrypt 1
p2p-detection debug-param protocol-param p2p_quic_ietf_frag_detection 1
```

Step 2 Enable P2P dynamic flow detection on the rulebase.

```
rulebase <rulebase_name>
p2p dynamic-flow-detection
```

Step 3 Enable SCONE processing on the rulebase.

```
rulebase <rulebase_name>
scone-enabled
or
rulebase <rulebase_name>
scone-enabled all
```

Step 4 Configure SCONE matching in the ruledef.

```
ruledef <ruledef_name>
scone-detected
```

Step 5 Configure SCONE throughput advice in the charging action.

```
charging-action <charging_action_name>
scone-throughput-advice enabled
or
charging-action <charging_action_name>
scone-throughput-advice enabled ta-value <bps>
or
charging-action <charging_action_name>
scone-throughput-advice disabled
```

The UPF detects SCONE-capable flows during flow establishment and signals throughput advice using either the configured `ta-value` or the bitrate derived from the most restrictive value between the flow MBR and APN-AMBR, depending on the charging-action configuration. The throughput advice is applied to eligible downlink SCONE traffic.

SCONE monitoring and troubleshooting

Monitor SCONE flow status and troubleshoot configuration issues using the following EDR attributes and show commands.

EDR attributes

Configure the following EDR attributes to monitor SCONE flow status and throughput advice information.

- **scone-enabled:** Indicates whether the flow was identified as SCONE-capable. Possible values are **y** and **n**.
- **scone-rate-signal:** Reports the last rate signal value applied to the flow. Possible values include a numeric rate signal, **bypass**, **ta-disable**, and **n/a**.

Show commands

Use the following commands to verify configuration and monitor SCONE statistics:

Command	Recommended Purpose
`show user-plane-service statistics all`	Displays node-level SCONE statistics.
`show subs user-plane-only flows`	Displays subscriber flow information, including SCONE status and Rate Signal.
`show subscribers user-plane-only full all`	Displays number of SCONE flows in the detailed session information for all subsribers.
`show subscribers user-plane-only callid <call-id>`	Displays number of SCONE flows.
`show subscribers user-plane-only callid <call-id> flow-id <flow-id>`	Displays SCONE information for the specified subscriber flow.
`show user-plane-service rulebase name <name>`	Displays SCONE configuration for the specified rulebase.
`show user-plane-service ruledef name <name>`	Displays SCONE configuration for the specified ruledef.

Command	Recommended Purpose
`show user-plane-service charging-action name <name>`	Displays SCONE throughput advice configuration for the specified charging action.

App-based Tethering Detection

Feature Description

The App-based Tethering Detection solution is built around the existing ADC plugins for App identifications. Tethering-specific patterns are added on top of recognized App plugins. These plugins successively return if the App flow is tethered or not. The App based Tethering Detection interworks with other existing supported tethering technique.



Note App-based Tethering Detection feature supports detection of only Netflix and YouTube applications.

Limitation

If there are any new TLS patterns used by tethered devices in the network, then those are not identified for tethering detection.

Configuring App-based Tethering Detection

This section describes how to enable support for App-based Tethering Detection.

Enabling App-based Tethering Detection at Rulebase Level

Use the following commands to enable App-based Tethering Detection for ADC traffic under ACS Rulebase Configuration Mode:

```
configure
  active-charging service service_name
    rulebase rulebase_name
      tethering-detection application
      exit
    exit
  exit
```

NOTES:

- If previously configured, use the **no tethering-detection** command to remove the tethering detection configuration from the rulebase.

App-based Tethering Detection at Ruledef Level

The following configuration allows you to define rule expressions to match tethered or non-tethered flows..

```
configure
  active-charging service service_name
    ruledef ruledef_name
     tethering-detection application { flow-tethered | flow-not-tethered
```

```
}
      exit
   exit
 exit
```

NOTES:

- If previously configured, use the **no tethering-detection** command to remove the tethering detection configuration from the ruledef.

Monitoring and Troubleshooting the App-based Tethering Detection

This section provides information regarding commands available to monitor and troubleshoot the App-based Tethering Detection.

Show Commands and Outputs

This section provides information regarding show commands and their outputs in support of this feature.

show user-plane-service statistic tethering-detection

The output of this CLI command has been enhanced to include the following fields in support of this feature.

- Tethering Detection Statistics (Application):
  - Total flows scanned
  - Tethered flows detected
  - Tethered uplink packets
  - Tethered downlink packets

show user-plane-service statistic rulebase name <rulebase_name>

The output of this CLI command has been enhanced to include the following fields in support of this feature.

- Tethering Detection (Application):
  - Total flows scanned
  - Tethered flows detected
  - Tethered uplink packets
  - Tethered downlink packets

Content Filtering

Content Filtering is an in-line service available for 3GPP and 3GPP2 networks to filter HTTP requests from mobile subscribers based on the URLs in the requests. This enables operators to filter and control the content that an individual subscriber can access, so that subscribers are inadvertently not exposed to universally unacceptable content and/or content inappropriate as per the subscribers' preferences.

Content Filtering Configuration

Use the following additional configuration to enable the content filtering:

```
configure
  require user-plane content-filtering
    content-filtering category database directory path path_address
    content-filtering category database max-version version_number
    end
```



Note The above configuration must be configured on the UPF, during boot time, to enable Content Filtering. Defining the above configuration post the User Plane configuration will lead to errors and inconsistencies.

Show Commands Input and/or Outputs

show subscribers user-plane-only callid *call_id* full all

SMF provides Content Filtering Policy ID in the Session Establishment/Modification Request. The following fields are displayed in support of this feature:

- SUBSCRIBER PARAMS
  - Content Filtering Policy ID

DNS Snooping

Charging

The charging of DNS Snooping takes place at SM-P.

Rule Definitions

Use the following CLI commands for specifying the rule definition hostnames (domain-names) and part of the host names.

```
ruledef <ruledef_name>
      ip [server-domain-name {contains|=|ends-with|starts-with} <url_string>]

      ip [server-domain-name {contains|=|ends-with|starts-with} <url_string>]

       multi-line-OR enabled
```

Use the no version of this CLI to delete the ruleline for ip server- domain-name.

```
ruledef <ruledef_name>
  no ip [server-domain-name {contains|=|ends-with|starts-with} <url_string>]

   exit
```

Use the following CLI for configurable timer of DNS entries at ECS level.

```
configure
        active-charging service service_name
```

```
ip dns-resolved-entries timeout <value_secs>
end
```

Whenever the ruledef containing the ip server-domain-name keyword is defined and used in rulebase, the ip-table is created per rulebase per instance.

Rule Matching

The functionality remains the same as the non-CUPS architecture.

Show CLIs

Use the following CLIs to check the table for DNS IP entries:**show user-plane-service [statistics dns-learnt-ip-addresses {summary | sessmgr instance <id> |all [verbose] }]**

Bulkstats

The following bulkstats are available in support of DNS Snooping feature:

- ecs-dns-learnt-ipv4-entries
- ecs-dns-flushed-ipv4-entries
- ecs-dns-replaced-ipv4-entries
- ecs-dns-overflown-ipv4-entries
- ecs-dns-learnt-ipv6-entries
- ecs-dns-flushed-ipv6-entries
- ecs-dns-replaced-ipv6-entries
- ecs-dns-overflown-ipv6-entries

The above bulkstats are added in the ECS schema same as in the non-CUPS architecture.

Note The SNMP Trap generation commands are not supported in CUPS DNS snooping feature.

Event Data Records

Table 36: Feature History

Feature Name	Release Information	Feature Description
S-NSSAI Support in EDR	2024.02.0	UPF supports a new EDR attribute **sn-nssai** as part of Flow and Transaction EDRs in the EDR Format configuration. This attribute reports the S-NSSAI value that UPF receives from SMF in PFCP Session Establishment Request.

Feature Name	Release Information	Feature Description
EDR Last Uplink and Downlink Packet Time Attributes	2023.04	The EDR configuration supports two new attributes **sn-last-uplink-pkt-time** and **sn-last-downlink-pkt-time** to identify the data stall issue.

Feature Description

Event Data Records (EDR) are usage records with support to configure content information, format, and generation triggers by the system administrative user.

When a flow is terminated, the UPF generates EDRs with detail information of the terminated flow.

How It Works

EDRs are generated from the User Plane on flow termination. During call setup and call modification, all call-specific attributes that are required for EDR generation is sent from SMF to UPF as part of the "Subscriber Params" IE within the Sx Establishment/Modification request messages.

On flow termination, the charging counters are fetched from VPP. All configured call-level attributes in the EDR format configuration along with the charging/volume counter attributes is sent to the CDRMOD proclet. This proclet writes these records to a file/disk, which is transferred to a configured external server.

TCP Fast Open

TCP Fast Open (TFO) is an extension to speed up the opening of successive TCP connections between two endpoints. It works by using a TFO cookie (a TCP option), which is a cryptographic cookie that is stored on the client and set upon the initial connection with the server. When the client reconnects, it sends the initial SYN packet along with the TFO cookie data to authenticate itself. If successful, the server starts sending data to the client even before the reception of the final ACK packet of the three-way handshake. Due to this, the difference between the following packets are recorded to calculate and record the time difference between control packets of TCP flow in EDR:

- SYN and SYN-ACK packet
- SYN-ACK and ACK packet

For information about rule variables that are added to capture the information in EDRs, refer *Configuring Additional TCP Fields* section.

Transaction Complete EDR

Transaction Complete EDRs are generated for HTTP EDRs when an HTTP transaction is completed. On completion, the charging counters are fetched from VPP. All configured call-level attributes in the EDR format configuration along with the charging/volume counter attributes is sent to the CDRMOD proclet. This proclet writes these records to a file/disk, which is transferred to a configured external server.

The following EDR attributes are supported:

- sn-start-time
- sn-end-time
- sn-start-time format MM/DD/YYYY-HH:MM:SS:sss

- sn-end-time format MM/DD/YYYY-HH:MM:SS:sss
- radius-calling-station-id
- radius-called-station-id
- bearer 3gpp imsi
- bearer 3gpp imei
- bearer 3gpp rat-type
- bearer 3gpp user-location-information
- ip subscriber-ip-address
- ip dst-address
- sn-ruledef-name
- sn-subscriber-port
- sn-server-port
- sn-app-protocol
- sn-volume-amt ip bytes uplink
- sn-volume-amt ip bytes downlink
- sn-flow-start-time format seconds
- sn-flow-end-time format seconds
- sn-volume-amt ip pkts uplink
- sn-volume-amt ip pkts downlink
- sn-direction
- traffic-type
- p2p protocol
- p2p app-identifier tls-cname
- p2p app-identifier tls-sni
- p2p app-identifier quic-sni
- bearer 3gpp sgsn-address
- sn-rulebase
- sn-charging-action
- flow tethered-ip-ttl
- flow ttl
- flow ip-control-param
- bearer qci

- tcp flag
- ip server-ip-address
- sn-flow-id
- sn-closure-reason
- sn-duration
- ip src-address
- ip protocol
- sn-charge-volume ip bytes uplink
- sn-charge-volume ip bytes downlink
- sn-nssai
- sn-dscp-uplink
- sn-dscp-downlink

The following HTTP EDR attributes are supported:

- http url length 2000
- http request method
- http content type
- http user-agent length 255
- http reply code
- http referer
- http host
- http cookie
- http header-length
- transaction-uplink-bytes
- transaction-downlink-bytes

The following DNS EDR attributes are supported:

- dns answer-ip-list
- dns answer-name
- dns previous-state
- dns query-name
- dns query-type
- dns return-code

- dns state
- dns tid

EDR Attributes to Identify Data Stalling

EDR supports the following attributes to identify the data stalling issue:

- **sn-last-uplink-pkt-time**
- **sn-last-downlink-pkt-time**

When the average difference is higher between the last packet on uplink or downlink direction and **sn-end time** (flow release time), it indicates data stalling.

These EDR variables do not capture the packet time of the control packets (FIN/RST/FIN-ACK/TCP TEAR DOWN ACK). The packet time for these packets is captured at the session manager using **sn-end time**.

For the offloaded flows, the attributes **sn-last-uplink-pkt-time** and **sn-last-downlink-pkt-time** is updated for packets going through the fastpath.

Limitations

The EDR feature in UPF has the following limitations:

- EDR will be generated only for flow end condition: Idle timeout, HAGR, normal flow termination, and during the end of a session.
- Charging-Action based EDR configuration is not supported.
- Reporting EDRs are not supported.
- The packet time of the control packet will not be updated for last_uplink_packet_time and last_downlink_packet_time except mid session TCP ACKs in EDR.

Configuring Event Data Records

Configuring EDRs on UPF

Use the following configuration to configure EDRs on UPF:

```
active-charging service service_name
   rulebase rulebase_name
      flow end-condition { timeout | normal-end-signaling | session-end |
 hagr } charging-edr charging_edr_format_name
     edr transaction-complete { http | dns } charging-edr
charging_edr_format_name
     exit
     edr-format format_name
        attribute attribute_name format time_format priority priority_val
```

NOTES:

- **flow end-condition**—This command allows you to configure the end condition of the session flows related to a user session and triggers EDR generation.
- **timeout**—Creates an EDR with the specified EDR format whenever a flow ends due to a timeout condition.

- **normal-end-signaling**—Creates an EDR with the specified EDR format whenever the flow end is signaled normally. For example, detecting FIN and ACK for a TCP flow, and create an EDR for the flow using the specified EDR format.
- **session-end**—Creates an EDR with the specified EDR format whenever a subscriber session ends. By this option, the session manager creates an EDR with the specified format name for every flow that has had any activity since the last EDR was created for the flow on session end.
- **charging-edr** *charging_edr_format_name*—Specifies the charging EDR format.
- **hagr**—Creates an EDR with the specified EDR format whenever a flow is terminated due to Inter-chassis Session Recovery action.
- **http**—Specifies HTTP protocol-related configuration.
- **dns**—Specifies DNS protocol-related configuration.
- **edr-format** *format_name*—Specifies an EDR that has to be configured.
- **attribute** *attribute_name* **format** *time_format* **priority** *priority_val*—Specifies the specific attribute to capture in EDR. **format** *time_format* specifies the timestamp format. It supports the following time formats:
  - MM/DD/YY-HH:MM:SS
  - MM/DD/YY-HH:MM:SS:sss
  - MM/DD/YYYY-HH:MM:SS
  - MM/DD/YYYY-HH:MM:SS:sss
  - YYYY/MM/DD-HH:MM:SS
  - YYYY/MM/DD-HH:MM:SS:sss
  - YYYYMMDDHHMMSS
  - YYYYMMDDHHMMSSsss
  - Seconds

  priority *priority_val* specifies specifies the position priority of the value within the EDR record. It must be an integer from 1 through 65535.

Configuration to Enable EDR Module

Use the following configuration to enable EDR module.

```
configure
  context context_name
    edr-module active-charging-service
    end
```

Configuring Additional TCP Fields

Prior to using the following CLI commands to configure additional TCP fields in the EDR, ensure that all the other EDR configurations are present.

```
configure
  active-charging service service_name
    edr-format edr_format_name
      rule-variable tcp syn-synack-rtt priority priority_value
      rule-variable tcp synack-ack-rtt priority priority_value
      end
```

Enabling EDR Billing

Use the following configuration for the system to generate an EDR immediately when a packet is received and it matches a ruledef that is associated with this charging action.

```
configure
  active-charging service service_name
    charging-action charging_action_name
      billing-action edr edr_format_name
      end
```

Configure the purging of pending DNS records

Follow these steps to configure to purging the pending DNS records.

Procedure

Step 1 Enter the Active Charging Service configuration mode.

```
configure
   active-charging service service_name
```

Step 2 Configure Maximum pending DNS records.

max-pending-dns-records *number_of_records*

The default value of **max-pending-dns-records** is 200.

Example:

```
[local]laas-si-setup#configure
[local]laas-si-setup(config)#active-charging service ACS
[local]laas-si-setup(config-acs)#max-pending-dns-records 890
```

Step 3 Verify the configuration.

Example:

```
[local]laas-si-setup# show configuration active-charging service name ACS verbose
config
  active-charging service ACS
    max-pending-dns-records 200


[local]laas-si-setup# show configuration active-charging service name ACS
config
  active-charging service ACS
```

```
#exit
max-pending-dns-records 890
```

Monitoring and Troubleshooting

show user-plane-service statistics rulebase name *rulebase_name*

The following fields are displayed in support of this feature:

- Rulebase Name
  - EDRs
  - Charge Volume
    - Uplink Pkts
    - Uplink Bytes
    - Downlink Pkts
    - Downlink Bytes
  - Charging EDRs
    - Total Charging EDRs generated
    - EDRs generated for handoff
    - EDRs generated for timeout
    - EDRs generated for normal-end-signaling
    - EDRs generated for session end
    - EDRs generated for rule match
    - EDRs generated for hagr
    - EDRs generated for flow-end content-filtering
    - EDRs generated for flow-end url-blacklisting
    - EDRs generated for content-filtering
    - EDRs generated for url-blacklisting
    - EDRs generated for any-error packets
    - EDRs generated for firewall deny rule match
    - EDRs generated for transaction completion
    - EDRs generated for voip call end
    - EDRs generated for dcca failure handling
    - EDRs generated for TCP optimization on
    - EDRs generated for tethering signature change

- EDRs generated for interim interval
- Total Flow-Overflow EDRs
- Total zero-byte EDRs suppressed

- Reporting EDRs
  - Total Charging REDRs generated
  - REDRs generated for handoff
  - REDRs generated for timeout
  - REDRs generated for normal-end-signaling
  - REDRs generated for session end
  - REDRs generated for rule match
  - REDRs generated for hagr
  - REDRs generated for flow-end content-filtering
  - REDRs generated for flow-end url-blacklisting
  - REDRs generated for content-filtering
  - REDRs generated for url-blacklisting
  - REDRs generated for any-error packets
  - REDRs generated for firewall deny rule match
  - REDRs generated for transaction completion
  - REDRs generated for voip call end
  - REDRs generated for dcca failure handling
  - REDRs generated for TCP optimization on
  - REDRs generated for tethering signature change

show user-plane-service edr-format all

The following fields are displayed in support of Additional TCP Fields in EDR feature:

- Service Name
  - EDR Format Name
    - rule-variable tcp syn-synack-rtt priority 1
    - rule-variable tcp synack-ack-rtt priority 2

Firewall Support

Table 37: Feature History

Feature Name	Release Information	Description
Firewall TCP Idle Timeout Action	2023.04	The Firewall feature inspects subscriber traffic performing IP session-based access control to protect subscribers from security attacks. UPF supports the TCP Idle Timeout action to drop the subscriber flow or send reset on TCP timeout expiry.

Feature Description

The Firewall feature in UPF 2.0 inspects subscriber traffic and performs IP session-based access control of individual subscriber sessions to protect the subscribers from malicious security attacks. UPF 2.0 supports the TCP Idle Timeout action to drop the subscriber flow or send reset on TCP timeout expiry.

The firewall configuration allows the system to inspect each packet of the subscriber data session. It also evaluates the security threat and applies the policies configured on uplink and downlink traffic. Firewall supports validation at per flow-level and per packet.

The Firewall feature supports the following functionality:

- Protection against DoS and DDoS attacks
- Application-level Gateway
- Stateful Packet Inspection and Filtering
- Stateless Packet Inspection and Filtering
- SNMP Thresholding

Subscriber Firewall Configuration

The firewall configurations are available under active charging configuration in UPF:

- Access-Ruledefs
- Firewall-NAT Policy

Enabling Default Firewall for IPv4 and IPv6

To configure the default firewall for IPv4 and IPv6, use the following configuration:

```
configure
  active-charging service service_name
    fw-and-nat policy policy_name
```

```
firewall policy ipv4-and-ipv6
end
```

Configuring TCP Idle Timeout Action

To configure the TCP Idle timeout action, use the following configuration:

```
configure
  active-charging service service_name
    fw-and-nat policy policy_name
      firewall tcp-idle-timeout-action { drop | reset }
      end
```

NOTES:

- **firewall tcp-idle-timeout-action { drop | reset }**—Specify the Stateful Firewall action to be taken on TCP idle timer expiry.
  - **drop**—Subscriber flow will be cleared or dropped without sending a reset on TCP timeout expiry.
  - **reset**—Specify to send a reset on TCP timeout expiry. This is the default value.
- The **firewall tcp-idle-timeout-action reset** CLI command is applicable only to firewall.
- UPF does not support flow mapping.

Along with the preceding service configuration, the following configuration is the default CLI behavior of various Firewall-related CLI within the service.

```
Dos-Protection:
      Source-Route                          : Disabled
      Win-Nuke                              : Disabled
      Mime-Flood                            : Disabled
      FTP-Bounce                            : Disabled
      IP-Unaligned-Timestamp                : Disabled
      Seq-Number-Prediction                 : Disabled
      TCP-Window-Containment                : Disabled
      Teardrop                              : Disabled
      UDP Flooding                          : Disabled
      ICMP Flooding                         : Disabled
      SYN Flooding                          : Disabled
      Port Scan                             : Disabled
      IPv6 Extension Headers Limit          : Disabled
      IPv6 Hop By Hop Options               : Disabled
      Hop By Hop Router Alert Option        : Disabled
      Hop By Hop Jumbo Payload Option       : Disabled
      Invalid Hop By Hop Options            : Disabled
      Unknown Hop By Hop Options            : Disabled
      IPv6 Destination Options              : Disabled
      Invalid Destination Options           : Disabled
      Unknown Destination Options           : Disabled
      IPv6 Nested Fragmentation             : Disabled

    Max-Packet-Size:
      ICMP                            : 65535
      Non-ICMP                        : 65535
    Flooding:
      ICMP limit                      : 1000
      UDP limit                       : 1000
      TCP-SYN limit                   : 1000
      Sampling Interval               : 1
```

```
TCP-SYN Flood Intercept:
  Mode                          : None
  Max-Attempts                  : 5
  Retrans-timeout               : 60
  Watch-timeout                 : 30
Mime-Flood Params:
  HTTP Header-Limit             : 16
  HTTP Max-Header-Field-Size    : 4096

No Firewall Ruledef Match Action:
  Uplink Action                 : permit
  Downlink Action               : deny

TCP RST Message Threshold     : Disabled
ICMP Dest-Unreachable Threshold    : Disabled
Action upon receiving TCP SYN packet with ECN/CWR Flag set    : Permit
Action upon receiving a malformed packet  : Deny
Action upon IP Reassembly Failure   : Deny
Action upon receiving an IP packet with invalid Options : Permit
Action upon receiving a TCP packet with invalid Options : Permit
Action upon receiving an ICMP packet with invalid Checksum: Deny
Action upon receiving a TCP packet with invalid Checksum: Deny
Action upon receiving an UDP packet with invalid Checksum: Deny
Action upon receiving an ICMP echo packet with id zero : Permit
TCP Stateful Checks : Enabled
First Packet Non-SYN Action: Drop
ICMP Stateful Checks: Enabled
TCP Partial Connection Timeout: 30
```

IP Reassembly

This section describes the reassembly functionality:

- In the non-CUPS architecture, fragments are buffered up to 64K bytes with the default Firewall configuration. All buffered and subsequent fragments are dropped beyond 64K bytes. In UPF, it is possible to reassemble the packet size to 9K bytes in a maximum of six fragments.
- The **firewall ip-reassembly-failure** CLI configures teardrop attack, nested fragmentation, and general ip-reassembly-failure. The maximum IP packet size is limited to six fragments (~9000 bytes).
- The following counters in firewall statistics get incremented for all attacks related to reassembly:
  - Downlink Packets Dropped due to IPv4 Reassembly Failure
  - Uplink Packets Dropped due to IPv4 Reassembly Failure
  - Downlink Dropped Bytes on IPv4 Reassembly Failure
  - Uplink Dropped Bytes on IPv4 Reassembly Failure
  - Downlink Packets Dropped due to IPv6 Reassembly Failure
  - Uplink Packets Dropped due to IPv6 Reassembly Failure
  - Downlink Dropped Bytes on IPv6 Reassembly Failure
  - Uplink Dropped Bytes on IPv6 Reassembly Failure

Monitoring and Troubleshooting

This section provides information about show commands and their output in support of this feature.

show config active-charging service name acs verbose

```
fw-and-nat policy SFW_NAT_TEST
      no firewall dos-protection source-router
      no firewall dos-protection winnuke
      no firewall dos-protection mime-flood
      no firewall dos-protection ftp-bounce
      no firewall dos-protection ip-unaligned-timestamp
      no firewall dos-protection tcp-window-containment
      no firewall dos-protection teardrop
      no firewall dos-protection flooding udp
      no firewall dos-protection flooding icmp
      no firewall dos-protection flooding tcp-syn
      no firewall dos-protection port-scan
      no firewall dos-protection ipv6-extension-hdrs
      no firewall dos-protection ipv6-hop-by-hop
      no firewall dos-protection ipv6-hop-by-hop router-alert
      no firewall dos-protection ipv6-hop-by-hop jumbo-payload
      no firewall dos-protection ipv6-hop-by-hop invalid-options
      no firewall dos-protection ipv6-hop-by-hop unknown-options
      no firewall dos-protection ipv6-dst-options
      no firewall dos-protection ipv6-dst-options invalid-options
      no firewall dos-protection ipv6-dst-options unknown-options
      no firewall dos-protection ipv6-frag-hdr nested-fragmentation
      no firewall dos-protection ip-sweep tcp-syn
      no firewall dos-protection ip-sweep udp
      no firewall dos-protection ip-sweep icmp
      firewall max-ip-packet-size 65535 protocol icmp
      firewall max-ip-packet-size 65535 protocol non-icmp
      firewall flooding protocol icmp packet limit 1000
      firewall flooding protocol udp packet limit 1000
      firewall flooding protocol tcp-syn packet limit 1000
      firewall flooding sampling-interval 1
      firewall tcp-syn-flood-intercept mode none
      firewall tcp-syn-flood-intercept watch-timeout 30
      firewall mime-flood http-headers-limit 16
      firewall mime-flood max-http-header-field-size 4096
      no firewall icmp-destination-unreachable-message-threshold
      access-rule no-ruledef-matches uplink action permit
      access-rule no-ruledef-matches downlink action deny
      firewall tcp-idle-timeout-action reset
      no firewall tcp-reset-message-threshold
      firewall tcp-syn-with-ecn-cwr permit
      firewall malformed-packets drop
      firewall ip-reassembly-failure drop
      no firewall validate-ip-options
      firewall tcp-options-error permit
      firewall icmp-echo-id-zero permit
      firewall icmp-checksum-error drop
      firewall tcp-checksum-error drop
      firewall udp-checksum-error drop
      firewall tcp-fsm first-packet-non-syn drop
      firewall icmp-fsm
      firewall policy ipv4-and-ipv6
      firewall tcp-partial-connection-timeout 30
      no nat policy
      no nat binding-record
      no nat pkts-drop edr-format
      no nat pkts-drop timeout
      default nat suppress-aaa-update
      nat private-ip-flow-timeout 180
      nat check-point-info basic limit-flows 100
      nat check-point-info sip-alg
      nat check-point-info h323-alg
```

```
  nat max-chunk-per-realm single-ip
#exit
```

Show CLIs

This section describes the supported show CLI commands for UPF:

- **show subscribers user-plane-only full all**
- **show subscribers user-plane-only flows**
- **show user-plane-service inline-services firewall statistics verbose**
- **show user-plane-service statistics rulebase all**
- **show alarm outstanding all**
- **show alarm outstanding all verbose**
- **show alarm statistics**
- **show user-plane-service statistics rulebase name** *rulebase_name*

SNMP Traps

This section describes the SNMP traps. To enable the trap, use the respective trap CLIs on the User Plane.

- **Dos-Attacks**—When the number of DoS attacks exceeds the set threshold value, the SNMP trap is generated. When the number falls below the threshold value within the configured time interval, the trap is cleared.
- **Drop-Packets**—When the number of packets dropped exceeds the threshold value, the SNMP trap is generated. When the number falls below the threshold value within the configured time interval, the trap is cleared.
- **Deny-Rule**—When the number of Deny Rules exceeds the threshold value, the SNMP trap is generated. When the number falls below the threshold value within the configured time interval, the trap is cleared.
- **No-Rule**—When the number of No Rules exceeds the threshold value, the SNMP trap is generated. When the number falls below the threshold value within the configured time interval, the trap is cleared.

Flow Idle Timeout Randomization

Every two seconds, the Session Manager polls the time of the latest packet from Session Manager instance, or the fastpath stream to determine idle flows. Short length flows become idle quickly as they are short due to the lesser number of packets and are short lived, within 5–10 seconds. As a result, large number of idle flows must be deleted due to the timeout at the given polling cycle of two seconds. Deletion of idle flows is CPU intensive as it involves statistics reconciliation, EDR generation, and fast path stream deletion. You can accomodate more flows with this feature as the short lived flows get cleared aggressively.

Configuring Flow Idle Timeout Randomization in ACS

Use the following configuration to randomize the idle timeout flow.

```
configure
active-charging service service_name
```

```
idle-timeout randomize-range range
  { default | no } idle-timeout randomize-range
    end
```

NOTES:

- **idle timeout**: Specifies the maximum duration that a flow can remain idle for, in seconds. Seconds must be an integer from 5 through 30. The flow will then be terminated based on the random value.
- **randomize-range**: Specifies the range of a period of time in seconds. The idle timeout applied, will be different for each flow.

  For example,

  `idle-timeout randomize-range 20`

  An integer random number is generated from 0 through 20. This number is added to the configured idle timeout value to check if the flow has become idle in the two second timer processing. If the idle timeout configured is 60 seconds, the actual timeout that is applied to each flow will be random in the range between *60 + 20* seconds causing staggered flow deletion.
- **no**: Disables the idle timeout randomization. This command is disabled by default.
- **default**: Configures the idle timeout randomization command with its default setting in seconds. Seconds must be an integer from 0 through 30. Default range is 0–30 seconds.

  For example, **default idle-timeout randomize-range** is equal to **idle-timeout randomize-range 30**.

HTTP URL Filtering

The HTTP URL Filtering feature simplifies rule definitions used for URL detection.

The HTTP request packet can have a proxy (prefixed) URL and an actual URL. If a proxy URL is found in the HTTP request packet, the HTTP URL Filtering feature truncates this URL from the parsed information and only the actual URL is used for rule matching and Event Data Records (EDR) generation.

Configuring the HTTP URL Filtering Feature

This section describes how to configure the HTTP URL Filtering feature.

Configuring Group of Prefixed URLs

To configure the group of prefixed URLs, use the following CLI commands:

```
configure
  active-charging service ecs_service_name
    group-of-prefixed-urls prefixed_urls_group_name
    end
```

Configuring URLs in the Group of Prefixed URLs

To configure URLs to be filtered in the group of prefixed URLs, use the following CLI commands:

```
configure
  active-charging service ecs_service_name
    group-of-prefixed-urls prefixed_urls_group_name
      prefixed-url url_1
      ...
```

```
        prefixed-url url_10
        end
```

Enabling the Group of Prefixed URLs in Rulebase

To enable the group of prefixed URLs in rulebase for processing prefixed URLs, use the following CLI commands:

```
configure
   active-charging service ecs_service_name
      rulebase rulebase_name
        url-preprocessing bypass group-of-prefixed-urls
prefixed_urls_group_name_1
        ...
        url-preprocessing bypass group-of-prefixed-urls
prefixed_urls_group_name_64
        end
```

This configuration on the control plane chassis will be pushed to the user plane with a PFD message for “group-of-prefixed-urls” and “rulebase-url-preprocessing”separately.

The group of prefixed URLs has the list of proxy URLs, which must be truncated. The rulebase contains multiple group of prefixed urls, which must be filtered. Charging ruledefs contain rules for actual URLs that must be searched after truncating URLs in the group of prefixed URLs.

Note

- Each group of prefixed URLs can have a maximum of ten prefixed URLs.
- A maximum of 64 group of prefixed URLs can be created and configured.

Show Commands

show user-plane-service group-of-prefixed-urls all | name *group_name*

This show command can be used on the user plane to verify whether the group of prefixed URLs are pushed or not. The output of this command is as follows:

- Name of the group of prefixed URLs
- Prefixed URLs
- Total number of prefixed URLs found

show user-plane-service rulebase name *rbase_name*

This show command can be used on the user plane to check whether the group of prefixed URLs is configured in rulebase or not. The output of this command is as follows:

- Name of rulebase
- Name of the groups of prefixed Urls for URL pre-processing

show user-plane-service statistics analyzer name http

The output of this command is as follows:

- Total HTTP Sessions
- Current HTTP Sessions
- Total Uplink Bytes
- Total Downlink Bytes
- Total Uplink Pkts
- Total Downlink Pkts
- Uplink Bytes Retrans
- Downlink Bytes Retrans
- Uplink Pkts Retrans
- Downlink Pkts Retrans
- Total Request Succeed
- Total Request Failed
- GET Requests
- POST Requests
- CONNECT Requests
- PUT requests
- HEAD requests
- Websocket Flows
- Invalid packets
- Wrong FSM packets
- Unknown request method
- Pipeline overflow requests
- Corrupt request packets
- Corrupt response packets
- Unhandled request packets
- Unhandled response packets
- Partial HTTP Header Anomaly prevented
- New requests on closed connection
- Memory allocation failures
- Packets after permanent failure
- Prefixed Urls Bypassed
- FastPath Statistics

- Total FP Flows
- Uplink (Total FP Pkts)
- Downlink (Total FP Pkts)
- Uplink (Total FP Bytes)
- Downlink (Total FP Bytes)



Note Prefixed URLs Bypassed counter has been added in http analyzer stats as a performance measurement to show the number of truncated prefixed URLs.

IP Readdressing

The IP Readdressing feature enables redirecting unknown gateway traffic based on the destination IP address of the packets to known/trusted gateways.

IP Readdressing is configured in the flow action defined in a charging action. IP readdressing works for traffic that matches a particular ruledef, and hence the charging action. IP readdressing is applicable to both uplink and downlink traffic. In the Enhanced Charging Subsystem, uplink packets are modified after packet inspection, rule matching, and so on, where the destination IP or port is determined, and replaced with the readdress IP or port just before they are sent out. Downlink packets (containing the readdressed IP or port) are modified when they are received, before the packet inspection, where the source IP or port is replaced with the original server IP or port number.

For one flow from an MS, if one packet is re-addressed, then all the packets in that flow is re-addressed to the same server. Features like DPI and rule-matching remain unaffected. Each IP address and port combination are defined as a ruledef.

In case of IP fragmentation, packets with successful IP reassembly are readdressed. However, IP fragmentation failure packets are not readdressed.

There are two different approaches for the readdress server selection, in case, server-list is configured under charging-action.

- **round-robin**: In round-robin approach, server selection happens in round-robin manner for every new flow. In round-robin, only active servers in the list are considered for selection.
- **hierarchy-based approach**: In hierarchy-based approach, servers are tagged as primary, secondary, tertiary, and so on, depending on the order they are defined in the readdress server-list. All flows are readdressed to the primary server until it is up and running. If Primary server goes down, then flows are readdressed to secondary server and the same logic goes on. Once primary server is active then flows switch back to primary server for readdressing.

An extra CLI is provided that enables user to select from hierarchy or round-robin approach for server selection. Both round-robin and hierarchy-based server selection approaches are applicable for both IPv4 and IPv6 based servers.

Configuring IP Readdressing

Readdressing of packets based on the destination IP address of the packets enables redirecting unknown gateway traffic to known/trusted gateways. This is implemented by configuring the re-address server in the charging action.

To configure the IP Readdressing feature, use the following configuration:

```
configure
  active-charging service acs_service_name
    charging-action charging_action_name
      flow action readdress server ipv4_address/ipv6_address [
discard-on-failure ]  [ dns-proxy-bypass ]  [ port port_number [
discard-on-failure ] [ dns-proxy-bypass ] ]
      end
```

To configure the IP Readdressing feature when the readdress server-list is defined under charging-action, use the following configuration:

```
configure
  active-charging service acs_service_name
    charging-action charging_action_name
      flow action readdress server-list server_list_name [ hierarchy ] [
round-robin ] [ dns-proxy-bypass ] [ discard-on-failure ]
      end
```

Following is the sample server-list configuration:

```
readdress-server-list DRE
consecutive-failures 1
response-timeout 10000
server 209.165.200.225 port 53
server 209.165.200.226 port 53
server 2001:420:54fe::1019 port 53
server 2001:420:54fe::1039 port 53
server 2001:420:54fe::1049 port 53
server 2001:420:54fe::1059 port 53
server 209.165.201.30 port 8080
#exit
```

Note A maximum of 10 servers can be configured in the list and a maximum of 10 lists can be configured in active-charging service.

Show Commands

This section provides information about the show CLI commands available in support of IP Readdressing feature.

CLI Command	Description
show user-plane-service readdress-server-list statistics all **show user-plane-service readdress-server-list statistics instance** *instance* **show user-plane-service readdress-server-list statistics name** *name*	Use these show CLI commands to display the readdress server list statistics.
show user-plane-service readdress-server-list statistics	

Use the **clear user-plane-service readdress-server-list statistics all**CLI to clear the readdress server list statistics.

Use the **show user-plane-service statistics charging-action all** CLI to check the "flows readdressed" counter.

L7 Protocol

The following L7 protocols are supported as part of DPI:

- DNS
- FTP
- HTTP
- HTTPS
- RTP/RTSP
- SIP

DNS

The UPF supports DNS protocol as part of L7 Analyzer.

FTP

The UPF 2.0 supports FTP proocol as part of L7 Analyzer.

HTTP

On completion of HTTP Request or HTTP Response, the uplink or downlink data packets are offloaded to VPP in the following cases:

- Content-Length—Volume-based offloading is supported for methods such as GET and POST. The HTTP flow with chunk-encoding data transfer mechanism does not get offloaded irrespective of the method defined in HTTP. If the stream is offloaded based on content-length, then the stream on the other end will also get offloaded until the former is not onloaded.
- CONNECT Method—The method where both uplink and downlink streams are offloaded after the flow is upgraded to CONNECT.

- WebSocket Method—After the flow is classified as WebSocket protocol, both uplink and downlink streams are offloaded.
- The streams are onloaded back in either of the following cases:
  - FIN packet received
  - Content-length is breached
  - PDN update

Header Parsing

Only the header fields defined in ruledefs, which are included in rulebase, are parsed. For x-header features, redirection is configured with dependencies on some HTTP header fields.

HTTP Charging

- Complete packets are charged.
- Partial packets are charged on completion. Packet completing the partial packet is also charged.
- Concatenated packets are charged.
- Delay Charging is enabled – Control packets are charged against application-based rule, depending on delay charging CLI configuration.
- Response-based charging is enabled – After HTTP request's response is received, then the HTTP request is charged against response rule's CA.

X-Header Parsing and Rule-Matching

Ruledefs with x-header rule-lines are parsed and matched.

WebSocket

Involves charging of subsequent packets of the flow after HTTP GET request as per the HTTP request, if the HTTP flow is upgraded to be a websocket flow.

Response-based TRM

Transactional Rule Matching is engaged after HTTP response packet is received.

X-Header Insertion

X-header Insertion is supported in HTTP Requests.

URL-based redirection

URL-based redirection is a UPF function that:

- returns a redirect response to the subscriber and terminates the existing TCP connections,
- triggers the subscriber's web browser to automatically send the original HTTP packet to a specified destination URL, and

- allows for the extraction of specific destination URLs from composite strings (e.g., returnURL) received via dynamic policies.

If the traffic flow is not HTTP, the UPF ignores the redirect URL option and forwards the packet normally. Redirection is supported only for specific HTTP packets (e.g., GET requests) sent in the uplink direction.

Table 38: Feature History

Feature Name	Release Information	Description
Redirect URL Encryption	2024.01	UPF supports the URL encryption of the dynamic fields using the Blowfish and AES encryption methods. Default Setting: Disabled – Configuration Required
Redirect Rewrite and TCP FIN/RST	2026.01.1	The ADC dynamic rule on the N7 interface performs Redirect URL Support. This is supported only for AES encryption.

Dynamic Fields and Encryption:

- **Field Insertion:** The UPF supports appending dynamic subscriber data to the redirected URL.
- **Security:** Sensitive dynamic fields can be protected using symmetric block ciphers.
  - **Blowfish:** Uses a 128-bit key for encryption and decryption.
  - **AES:** Uses Cipher Block Chaining (CBC) mode with 256-bit keys. The UPF uses a hardcoded value of 10 rounds for AES encryption.

Note Redirect URL on the N7 interface is supported only for AES encryption.

Restrictions for URL redirection using N7 interface

- This feature supports UL L7 packets only. It does not support DL L7 packets, as the feature requirement is specifically to support UL L7.
- Do not modify existing charging action configurations.

Supported dynamic fields for redirection

The UPF supports the following dynamic fields for insertion into redirected URLs. Fields marked with an asterisk (*) support encryption.

- #ACSMGR_BEARER_CALLED_STATION_ID#
- #BEARER.IMSI# *

- #BEARER.MSISDN# *
- #HTTP.HOST#
- #HTTP.URL#
- #HTTP.URI#
- #RTSP.URI#
- #RULEBASE#

The UPF supports the following dynamic fields for insertion into redirected URLs on N7 interface

- #BEARER.IMSI# *
- #HTTP.URL#

How UPF redirection and termination work

The UPF uses the service-scheme framework to detect specific policy events. It then executes redirection or flow termination based on ADC dynamic rules received over the N4 interface from the SMF.

Summary

The key components involved in the process are:

- PCF: provides dynamic policies containing composite redirect URLs or block instructions.
- SMF: relays policy information to the UPF over the N4 interface.
- UPF: Performs redirect and encryption, inserts dynamic data, and manages TCP signaling (FIN/RST) for blocked traffic.

Workflow

The workflow to enable redirect or termination by N7 Interface:

1. **service-scheme Triggering:** The service-scheme framework detects a far-received or pdr-received event during session establishment or modification.
2. **Condition Matching:** The system evaluates trigger-conditions, such as checking if a redirect URL contains specific keywords like returnURL or other keyword based on the configuration.
3. **Action Execution:** The UPF executes the mapped trigger-action, which points to a specific charging-action based on the adc rules coming during establishment or modification of session.
4. **URL Processing:** For redirection, the UPF extracts the actual URL, populates dynamic fields, and applies encryption.
5. **Flow Management:** For termination, the UPF sends a TCP FIN to the web client. It also sends a TCP RST to the web server to gracefully close the connection.

Configuring URL Encryption

To enable encryption of the dynamic fields in the redirected URL, use the following sample configuration:

```
configure
  active-charging service service_name
    charging-action charging_action_name
      flow action redirect-url redirect_url [ encryption { blowfish128 |
aes256 } [ encrypted ] key key ]
    end
```

NOTES:

- **flow action redirect-url** *redirect_url*—Specify the redirect URL. It can include up to 16 dynamic fields. *redirect_url* must be an alphanumeric string of 1 through 511 characters.
- **encryption { blowfish128 | aes256 } [encrypted] key** *key*—Specify to enable encryption for dynamic fields of the redirect URL.
  - **blowfish128**—Specify to use Blowfish encryption with 128-bit key for encrypting the dynamic fields.
  - **aes256**—Specify to use AES-CBC encryption with 256-bit key for encrypting the dynamic fields.
  - **encrypted**—Specify to encrypt the key.
  - **key** *key*—Specify the key to use for encryption of dynamic fields. *key* must be an alphanumeric string of 1 through 523 characters.
- If encoding is enabled in the URL to be redirected, the encrypted URL supports up to 2300 characters. If the length exceeds, then the URL will be truncated at 2300 characters.

Configure N7 based redirect flow action

Set up the UPF to process dynamic redirect URLs using N7 interface and apply encryption.

Before you begin

- Verify that the L7 Analyzer is active (required for HTTPS support).
- If using AES encryption, prepare a key phrase (alphanumeric, 1 to 523 characters).
- Ensure the redirect URL string is between 1 and 511 characters.

Follow these steps to configure redirect flow actions:

Procedure

Step 1 Define the ruledefs for URL identification.

Example:

```
[upf] # ruledef rule_name
 [upf] #  http url contains test.html
 [upf] #  http url !contains abc.html
 [upf] #  multi-line-or all-lines
```

Step 2 Configure the trigger-conditions for redirect URLs.

Example:

```
[upf] # trigger-condition TC1
 [upf] #  redirect-url contains "returnURL="
 [upf] #  redirect-url starts-with "https://snap"
```

Step 3 Define trigger-actions for redirection.

Example:

```
[upf] # trigger-action TA1
    [upf] #  execute-charging-action url_redirect_ca
```

Step 4 Create charging-actions with encryption parameters.

Example:

```
[upf] # charging-action url_redirect_ca
[upf] #  flow action redirect-url <return url recieved over n4 interface with ADC rule>
```

Step 5 Bind the trigger to the service-scheme.

Example:

```
[upf] # service-scheme SS1
 [upf] #  trigger far-received
   [upf] #  priority 1 trigger-condition TC1 trigger-action TA1
```

Configure terminate flow action

Implement traffic blocking using TCP signaling.

Before you begin

Follow these steps to configure flow termination:

Procedure

Step 1 Define the ruledefs for traffic blocking.

Example:

```
[upf] # ruledef terminate_rule
  [upf] #  ip any-match = TRUE
```

Step 2 Configure the trigger-conditions for termination.

Example:

```
[upf] # trigger-condition TC2
  [upf] # tdf-appid contains <terminate app id recived with ADC dynamic rule over N4 interface>
```

Step 3 Define trigger-actions for termination.

Example:

```
[upf] # trigger-action TA2
   [upf] #   execute-charging-action terminate_all
```

Step 4 Create the charging-action for flow termination.

Example:

```
[upf] # charging-action terminate_all
   [upf] #  flow action terminate-flow
```

Step 5 Bind the trigger to the service-scheme.

Example:

```
[upf] # service-scheme SS1
[upf] #   trigger pdr-received
[upf] #   priority 2 trigger-condition TC2 trigger-action TA2
```

Redirect and terminate flow statistics

You can use the following commands to verify the feature status and view statistics for redirected or terminated flows.

Subscriber FAR Details

To view the extracted redirect URL and the applied charging action for a specific call, use the following syntax:

Example:

```
[upf] # show subscribers user-plane-only callid [call_id] far full all
```

Sample Output:

```
Callid: 00004e29
Interface Type: Sxb
 FAR-ID       : 0x0003
 Destination Interface : Core
 Apply Action     : FORWARD
 Outer Header Creation   :
 Remote TEID     : 0x0
 Transport Level Marking    : N/A
 Transport Level Marking Options :
   Copy Inner:                      No
   Copy Outer:                      No
 Inner Packet Marking            : N/A
 Layer2 Marking                  : N/A
 Remote IP Address   :
 Remote Port    : N/A
 Number of Associated PDRs: 1
 Linked Traffic Endpoint : 0x2
 Redirect Info  :
Redirect Address Type  : URL
Redirect Address Value :
https://snap.mobile.att.net/?nonce=ApolloHONDA&returnURL=https://myvehicle.att.com/%23/start%3Foem=honda

       Charging Action Name           : url_redirect_ca
```

Trigger Action Configuration

To verify which charging actions are mapped to specific trigger actions, use the following syntax:

Example:

```
[upf] # show user-plane-service trigger-action all
```

Sample Output:

```
Trigger-Action: TA1
       HTTP Response Based TRM      : none
       HTTP Response Based Charging : none
       Throttle Suppress            : Disabled
       Flow Recovery                : Disabled
       Traffic Optimization         : Disabled
       Step Up GBR                  : Disabled
       Step Down GBR                : Disabled
       TCP Acceleration             : Disabled
       TCP Acceleration Threshold   : Disabled
       Service-Chain                : none
       UP-Service-Chain             : none
       Flow-IP-Readdressing         : Disabled
       EDNS-Encode                  : Disabled
       Charging-Action-Map          : none
       Charging-Action-Name         : url_redirect_ca

Trigger-Action: TA2
       HTTP Response Based TRM      : none
       HTTP Response Based Charging : none
       Throttle Suppress            : Disabled
       Flow Recovery                : Disabled
       Traffic Optimization         : Disabled
       Step Up GBR                  : Disabled
       Step Down GBR                : Disabled
       TCP Acceleration             : Disabled
       TCP Acceleration Threshold   : Disabled
       Service-Chain                : none
       UP-Service-Chain             : none
       Flow-IP-Readdressing         : Disabled
       EDNS-Encode                  : Disabled
       Charging-Action-Map          : none
       Charging-Action-Name         : terminate_all
```

Trigger Condition Details

To review the criteria used to select charging actions, use the following syntax:

Example:

```
[upf] # show user-plane-service trigger-condition all
```

Sample Output:

```
Trigger-Condition: TC1
       redirect-url  contains returnURL=
       redirect-url  starts-with https://snap
       redirect-url  contains nonce=Name
       Multi-line-OR All lines : Disabled

Trigger-Condition: TC2
       Tdf-AppId  contains terminate_
       Multi-line-OR All lines : Disabled
```

Global Redirect Statistics

To view global counts for redirected flows, use the following syntax:

Example:

```
[upf] # show user-plane-service statistics all
```

Sample Output:

```
ADC Redirect Stats:
    ADC Redirected Flows         : 3
```

Terminated Flow Counts

To view the number of terminated flows for a specific call, use the following syntax:

Example:

```
[upf] # show subscribers user-plane-only full callid <call_id>
```

Sample Output:

```
Flow Action Terminated Flows: 5
```

PDR Termination Details

To view the number of packets matched while terminating flows and the applied charging action, use the following syntax:

Example:

```
[upf] # show subscribers user-plane-only callid <call_id> pdr full all
```

Sample Output:

```
Callid: 12345
Interface Type: Sxb
IP address: 11.11.0.1
PDR-ID:                    0x0003
Rule Name:                dynamicfacebook
Charging Action Name:      terminate_all
TDF App Id:   terminate_rule
TDF Notifications:   Enabled
Hits:                      5 Match Bypassed:            0
Matched Bytes:          285 Matched Packets:            5
```

HTTPS

The UPF supports HTTPS protocol as part of L7 Analyzer.

RTP/RTSP

The UPF 2.0 supports RTP and RTSP protocols as part of L7 Analyzer.

SIP

Session Initiation Protocol is an IETF-defined signaling protocol widely used for controlling communication sessions such as voice and video calls over Internet Protocol (IP). The protocol can be used for creating, modifying and terminating two-party (unicast) or multiparty (multicast) sessions. Sessions may consist of one or several media streams.

The UPF 2.0 supports SIP as part of L7 Analyzer.

Monitoring and Troubleshooting

This section provides information regarding the CLI command available in support of monitoring and troubleshooting the feature.

Show Command(s) and/or Outputs

This section provides information regarding show commands and/or their outputs in support of this feature.

DNS

Use the following CLI command to get statistics related to DNS:

```
show user-plane-service statistics analyzer name dns
```

FTP

Use the following CLI command to get statistics related to FTP:

```
show user-plane-service statistics analyzer name ftp
```

HTTP

Use the following CLI command to get statistics related to HTTP:

```
show user-plane-service statistics analyzer name http
```

HTTPS

Use the following CLI command to get statistics related to HTTPS:

```
show user-plane-service statistics analyzer name secure-http
```

RTP

Use the following CLI command to get statistics related to RTP:

```
show user-plane-service statistics analyzer name rtp
```

RTSP

Use the following CLI commands to get statistics related to RTSP:

- `show user-plane-service statistics analyzer name rtsp`
- `show user-plane-service statistics analyzer name rtsp verbose`

SIP

Use the following CLI command to get statistics related to SIP:

```
show user-plane-service statistics analyzer name sip
```

NAT Support on N4

Table 39: Feature History

Feature Name	Release Information	Description
NAT Support on the N4 interface	2023.04	The Network Address Translation (NAT) feature translates non-routable private IP addresses to routable public IP addresses from a pool of public IP addresses. UPF supports NAT on the N4 interface to configure network addresses and send NAT binding records to N4. The NAT policy and IP pool for NAT public IP addresses are configured on UPF for N4.

Feature Description

The Network Address Translation (NAT) feature translates non-routable private IP address(es) to routable public IP address(es) from a pool of public IP addresses. This feature enables the user to conserve on the number of public IP addresses required to communicate with external networks. This feature also ensures security as the IP address scheme for the internal network is masked from external hosts, and each outgoing and incoming packet goes through the translation process.

UPF supports NAT on the N4 interface to configure network addresses. The system is configured to automatically forward data packets encapsulating the source IP or source port address of the UE with NAT IP address and NAT port.

The supported NAT combinations include:

- NAT44 On Demand Many to One
- NAT44 On Demand One to One
- NAT64 On Demand Many to One
- NAT 64 On Demand One to One
- NAT44 Not On Demand Many to One
- NAT44 Not On Demand One to One
- NAT64 Not On Demand Many to One
- NAT64 Not On Demand One to One

For supplemental information about NAT, see the StarOS *NAT Administration Guide*.

NOTE: Not all features and/or functionality that are mentioned in the StarOS *NAT Administration Guide* are applicable to UPF.

Limitations

NAT support on the N4 interface has the following limitations:

- Supports only NAT44 with many-to-one and on-demand modes.
- All NAT pools are configured at respective User Plane in the destination context.
- Does not support charging action with CLI action deny in fw-and-nat policy and flow-any-error charging action in active-charging-service.
- Does not support access-rules which are configured with "dynamic-only" and "static-and-dynamic" rules from external servers.
- Does not support multiple IPs from the same realm.
- Does not support the next hop forwarding in NAT pool.
- Does not support the port range in NAT pool.
- Does not support the skip private IP check CLI.
- Does not support RADIUS and Gy returned Fw-and-nat policy-based applying NAT policy.
- Does not support the bearer-specific filters in access-ruledefs.
- Access-rules do not support trigger open-port port range configuration in fw-and-nat policy.
- Does not support the NAT port recovery (fw-and-nat action) after SR/ICSR.
- Does not support the NAT reassembly timeout CLI in active-charging service. The generic context-level CLI on UP must be used instead.
- Does not support the NAT fragmentation reassembly failure.
- Does not support the NAT flow-mapping timer.
- For N:M redundancy, the NAT IP pools to be configured from RCM done as part of interface configuration for each UP host and the pool name must be unique across all active user planes. It is mandatory to use NAT groups for all pools so that the same NAT realm referred in fw-and-nat policy can be applicable to all user planes.
- In N:M redundancy, the total number of NAT IP pools collectively configured on all UPs through RCM must be as per the maximum limit (2000) of IP pools. The configuration in standby User Plane fails if the cumulative total of all active UPs exceeds the maximum value.

Configuring NAT

The NAT policy (Firewall-and-NAT policy) and access-ruledef configurations are made under active-charging service in UPF.

For information on NAT-related CLI commands, refer to the *StarOS NAT Administration Guide > NAT Configuration* chapter.



Note Not all CLI commands and configurations mentioned in the *StarOS NAT Administration Guide > NAT Configuration* chapter are applicable to UPF.

Sample Configurations

User Plane

The following pool-related configuration is required at User Plane in ISP context.

```
configure
   context ISP1-UP
      ip pool NAT44_PUBLIC1 209.165.200.225 255.255.255.224 napt-users-per-ip-address 2
on-demand port-chunk-size 16 max-chunks-per-user 4 group-name NAT44_GRP1
      ip pool NAT44_PUBLIC2 209.165.200.226 255.255.255.224 napt-users-per-ip-address 2
on-demand port-chunk-size 16 max-chunks-per-user 4 group-name NAT44_GRP1
      ip pool NAT44_PUBLIC3 209.165.200.227 255.255.255.224 napt-users-per-ip-address 2
on-demand port-chunk-size 8 max-chunks-per-user 1 group-name NAT44_GRP2
      ip pool NAT44_PUBLIC4 209.165.200.228 255.255.255.224 napt-users-per-ip-address 4
on-demand port-chunk-size 32256 max-chunks-per-user 4 group-name NAT44_GRP2
      ip pool NAT44_PUBLIC5 209.165.200.229 255.255.255.224 napt-users-per-ip-address 8064
 on-demand port-chunk-size 8 max-chunks-per-user 2
      end
```

Sample NAT Pool Related Configuration for Different NAT Pool Types

```
1-1 on-demand:
---------------------------
config
context ISP1-UP
ip pool NAT44_ipv4_1_1 209.165.200.230 255.255.255.224 nat-one-to-one on-demand
nat-binding-timer 60
end

N-1 Not-on-demand:
------------------------
config
context ISP1-UP
ip pool NAT44_ipv4_N_1 209.165.200.231 255.255.255.224 napt-users-per-ip-address 2
max-chunks-per-user 2 port-chunk-size 8
end

1-1 Not-on-demand:
------------------------
config
context ISP1-UP
ip pool NAT44_ipv4_NOD_1_1 209.165.200.232 255.255.255.224 nat-one-to-one
end
```



Note Add the Control Plane configuration and map one or more access rule definitions to any required NAT Pool or Group configured in the User Plane.

Monitoring and Troubleshooting

Gathering NAT Statistics

The following table lists the commands that can be used to gather NAT statistics.

The first column lists the statistics to gather and the second column lists the command to use.

Statistics/Information	Show Command
Information for all current subscribers who have either active or dormant sessions. Checks IP address associated with subscriber. Also displays all the IP addresses that are in use in a NAT realm.	**show subscribers user-plane-only full all**
Information on NAT subsystem statistics.	**show user-plane-service statistics all**
All NAT-related statistics.	**show user-plane-service statistics nat all**
All NAT Realm-related statistics.	**show user-plane-service statistics nat nat-realm all**
Statistics of all NAT IP pools in a NAT IP pool group.	**show user-plane-service statistics nat nat-realm** *pool_name*
Information on NAT bind records generated.	**show user-plane-service edr-format statistics all**
Verifying association of fw-and-nat policy in APN on UP.	**show user-plane-service pdn-instance name** *name*
Verifying cofiguration of fw-an-nat policy on UP.	**show user-plane-service fw-and-nat policy all**
Information on NAT bind records generated for port chunk allocation and release.	**show user-plane-service rulebase name** *name*
Information on access ruledef.	**show user-plane-service ruledef all**
Verifying association of fw-and-nat policy in rulebase on UP.	**show user-plane-service rulebase name** *name*

Clear Commands

The following clear CLI commands are available in support of this feature:

- **clear user-plane-service statistics nat nat-realm all**
- **clear user-plane-service statistics nat all**

SNMP Traps for NAT Parameter Thresholds

The following SNMP traps for NAT parameter thresholds are supported.

SNMP Traps	Description
ThreshNATPortChunks	Generated when NAT port chunk usage reaches configured threshold limit
ThreshClearNATPortChunks	Generated when NAT port chunk usage reaches configured clear threshold limit.
ThreshNATPktDrop	Generated when NAT packet drop reaches configured threshold limit.
ThreshClearNATPktDrop	Generated when NAT packet drop reaches configured clear threshold limit.
ThreshIPPoolUsed	Generated when the number of IPs used in the IP Pool reaches configured threshold limit.

SNMP Traps	Description
ThreshClearIPPoolUsed	Generated when the number of IPs used in the IP Pool reaches configured clear threshold limit.
ThreshIPPoolFree	Generated when IP pool is free and threshold limit reached.
ThreshClearIPPoolFree	Generated when IP pool is used, and clear threshold limit reached.
ThreshIPPoolAvail	Generated when IP pool is available for next flow and configured threshold reached.
ThreshClearIPPoolAvail	Generated when IP pool is used, and configured threshold is reached.

NOTE: The respective CLIs must be configured in the User Plane to enable these traps.

Bulk Statistics

NAT-realm Schema

The NAT realms are configured in User Plane and statistics are stored per-context per-realm. These statistic variables, both cumulative and snapshot, are available in the nat-realm schema.

Table 40: NAT-realm Schema

Variable Name	Data Type	Counter Type	Description
Vpnname	String	Info	Context name.
Realmname	String	Info	Realm name.
nat-rlm-bind-updates	Int64	Counter	Total interim AAA NBU sent.
nat-rlm-bytes-txferred	Int64	Counter	Total number of NAT44 and NAT64 bytes transferred by realm (uplink + downlink).
nat-rlm-bytes-nat44-tx	Int64	Counter	Total number of NAT44 bytes transferred by realm.
nat-rlm-bytes-nat64-tx	Int64	Counter	Total number of NAT64 bytes transferred by realm.
nat-rlm-ip-flows	Int64	Counter	Total number of NAT44 and NAT64 flows used by the realm.
nat-rlm-nat44-flows	Int64	Counter	Total number of NAT44 flows processed by realm.
nat-rlm-nat64-flows	Int64	Counter	Total number of NAT64 flows processed by realm.
nat-rlm-ip-denied	Int32	Counter	Total number of NAT44 and NAT64 flows denied NAT IP address.
nat-rlm-ip-denied-nat44	Int64	Counter	Total number of NAT44 flows denied IP.

Variable Name	Data Type	Counter Type	Description
nat-rlm-ip-denied-nat64	Int64	Counter	Total number of NAT64 flows denied IP.
nat-rlm-port-denied	Int32	Counter	Total number of NAT44 and NAT64 flows denied ports.
nat-rlm-port-denied-nat44	Int64	Counter	Total number of NAT44 flows denied ports.
nat-rlm-port-denied-nat64	Int64	Counter	Total number of NAT64 flows denied ports.
nat-rlm-memory-denied	Int64	Counter	Total number of NAT44 and NAT64 flows denied memory.
nat-rlm-memory-denied-nat44	Int64	Counter	Total number of NAT44 flows denied memory.
nat-rlm-memory-denied-nat64	Int64	Counter	Total number of NAT64 flows denied memory.
nat-rlm-ttl-ips	Int32	Gauge	Total number of NAT public IP addresses, per context per NAT realm. Is a static value.
nat-rlm-ips-in-use	Int32	Gauge	Total number of NAT IP addresses currently in use, per context per NAT realm.
nat-rlm-current-users	Int32	Gauge	Total number of subscribers currently using the NAT realm.
nat-rlm-ttl-port-chunks	Int32	Gauge	Total number port-chunks, per context per NAT realm. Is a static value.
nat-rlm-chunks-in-use	Int32	Gauge	Total number of port-chunks currently in use, per context per NAT realm.
nat-rlm-port-chunk-size	Int32	Gauge	Size of the port chunk in the NAT realm.
nat-rlm-port-chunk-average-usage-tcp	Int32	Gauge	Average TCP port usage in the allocated TCP ports, i.e. out of allocated TCP ports how many got used. Not percentage value.
nat-rlm-port-chunk-average-usage-udp	Int32	Gauge	Average UDP port usage in the allocated UDP ports, i.e. out of allocated UDP ports how many got used. Not percentage value.
nat-rlm-port-chunk-average-usage-others	Int32	Gauge	Average other (ICMP or GRE) port usage in the allocated other ports, i.e. out of allocated ‘other’ ports how many got used. Not percentage value.
nat-rlm-max-port-chunk-subs	Int64	Counter	Total number of subscribers who used maximum number of port chunks.
nat-rlm-max-port-chunk-used	Int32	Counter	Maximum port chunks used.

Variable Name	Data Type	Counter Type	Description
nat-rlm-max-cur-port-chunk-subs	Int64	Gauge	Current number of subscribers using maximum number of port chunks.
nat-rlm-max-cur-port-chunk-used	Int32	Gauge	Maximum port chunks used by active subscribers.

NAT Binding Records

Whenever a NAT IP address or NAT port-chunk is allocated/deallocated to/from a subscriber, NAT Binding Records (NBR) can be generated. Generation of NBRs is configurable in the Firewall-and-NAT policy configuration.

Sample NBR

```
#sn-start-time,sn-end-time,ip-protocol,ip-subscriber-ip-address,ip-server-ip-address,sn-subscriber-port,
sn-server-port,sn-nat-ip,sn-nat-port-block-start,sn-nat-port-block-end,sn-subscriber-nat-flow-ip,sn-subscriber-nat-flow-port,
sn-nat-realm-name,sn-nat-subscribers-per-ip-address,sn-nat-binding-timer,sn-nat-gmt-offset,sn-nat-port-chunk-alloc-dealloc-flag,
sn-nat-port-chunk-alloc-time-gmt,sn-nat-port-chunk-dealloc-time-gmt,sn-nat-no-port-packet-dropped,sn-closure-reason
,,,209.165.200.225,,,,209.165.201.1,1024,1039,,,NAT44_PUBLIC2,2,60,+0530,1,02/18/2020
06:41:08,,,
,,,209.165.200.225,,,,209.165.201.2,1024,1031,,,NAT44_PUBLIC5,8064,60,+0530,1,02/18/2020
06:41:14,,,
,,,209.165.200.225,,,,209.165.201.3,1024,1039,,,NAT44_PUBLIC2,2,60,+0530,0,02/18/2020
06:41:08,02/18/2020 06:42:12,,
,,,209.165.200.225,,,,209.165.201.14,1024,1031,,,NAT44_PUBLIC5,8064,60,+0530,0,02/18/2020
06:41:14,02/18/2020 06:44:24,,
```

EDRs

The following NAT-specific attributes are supported in regular EDRs:

- sn-nat-subscribers-per-ip-address: Subscriber(s) per NAT IP address
- sn-subscriber-nat-flow-ip: NAT IP address of NAT-enabled subscribers
- sn-subscriber-nat-flow-port: NAT port number of NAT-enabled subscribers

Sample EDR

```
#sn-start-time,sn-end-time,ip-protocol,ip-subscriber-ip-address,ip-server-ip-address,sn-subscriber-port,sn-server-port,
sn-nat-ip,sn-nat-port-block-start,sn-nat-port-block-end,sn-subscriber-nat-flow-ip,sn-subscriber-nat-flow-port,sn-nat-realm-name,
sn-nat-subscribers-per-ip-address,sn-nat-binding-timer,sn-nat-gmt-offset,sn-nat-port-chunk-alloc-dealloc-flag,sn-nat-port-chunk-alloc-time-gmt,
sn-nat-port-chunk-dealloc-time-gmt,sn-nat-no-port-packet-dropped,sn-closure-reason
02/18/2020 12:11:11:630,02/18/2020
12:11:11:632,1,209.165.200.225,209.165.201.1,0,0,,,,209.165.200.230,1024,,2,,,,,,0,0
02/18/2020 12:11:08:672,02/18/2020
12:11:09:671,6,209.165.200.225,209.165.201.1,1001,3000,,,,209.165.200.230,1034,,2,,,,,,0,0
02/18/2020 12:11:14:499,02/18/2020
12:11:14:499,17,209.165.200.225,209.165.201.1,1001,3000,,,,209.165.200.240,1025,,8064,,,,,,0,0
```

Packet Drop EDR

Sample Packet Drop EDR

```
#sn-nat-no-port-packet-dropped,sn-start-time,sn-end-time,sn-subscriber-imsi
2,03/13/2020 08:28:24,03/13/2020 08:28:54,123456789012345
```

Tethering Detection

Feature Description

Tethering refers to the use of a mobile smartphone as a USB dongle or modem to provide Internet connectivity to PC devices (laptops, PDAs, tablets, and so on) running on the smartphone's data plan. Typically, for smartphone users, most operators have in place an unlimited data plan, the usage of which is intended to be from the smartphone as a mobile device. However, some subscribers use the low cost or unlimited usage data plan to provide Internet connectivity to their laptops in places where normal Internet connection through broadband or Wi-Fi may be costly, unavailable, or insecure.

The Tethering Detection feature enables detection of subscriber data traffic originating from PC devices tethered to mobile smartphones, and also provides effective reporting to enable service providers take business decisions on how to manage such usage and to bill subscribers accordingly. Tethering Detection is supported for IPv4 (TCP) and IPv6 traffic flows.

In this release, IP-TTL based tethering is supported. This feature is configurable at the rulebase level and is applicable on all flows for all subscribers having IP-TTL configuration within the rulebase.

Configuring Tethering Support

This section describes how to configure the Tethering Support feature.

Configuring the Tethering Support feature involves the following steps:

- Rulebase Configuration for Tethering
- Ruledef Configuration for Tethering
- EDR Configuration for Tethering

Rulebase Configuration for Tethering

Use the following commands to configure the rulebase parameters for tethering.

```
configure
  active-charging service service_name
    rulebase rulebase_name
      tethering-detection ip-ttl valuettl_value
      end
```

NOTES:

- **tethering-detection**: This command allows you to enable/disable the Tethering Detection feature for the current rulebase, and specifies the database to use.
- **ip-ttl value** *ttl_value*: Specifies to perform tethering detection using IP-TTL configuration. *ttl_value* must be an integer from 1 through 255 to configure TTL values for tethered flows.

Ruledef Configuration for Tethering

Use the following commands to configure ruldef parameters for tethering.

```
configure
   active-charging service service_name
      ruledef ruledef_name
         ip any-match operator_condition
         tethering-detection ip-ttl flow-tethered
         end
```

NOTES:

- **ip any-match** *operator_condition*: This command allows you to define rule expressions to match all IPv4/IPv6 packets.
- **ip-ttl**: Specifies to select flows that were tethered or non-tethered as per IP-TTL values.
- **flow-tethered**: Specifies to match if tethering is detected on flow.

EDR Configuration for Tethering

Use the following commands to configure EDR for tethering:

```
configure
   active-charging service service_name
      edr-format format_name
         rule-variable flow tethered-ip-ttl priority priority_value
         rule-variable flow ttl priority priority_value
         end
```

NOTES:

- **edr-format** *format_name*: configures EDR formats.
- **flow**: Configures the flow related fileds in an EDR.
- **tethered-ip-ttl**: IP-TTL based tethering detected on flow.
- **ttl**: Time To Live/Max hops value received in the first packet of the flow.

Monitoring and Troubleshooting

This section provides information regarding the CLI command available in support of monitoring and troubleshooting the feature.

Show Command(s) and/or Outputs

This section provides information regarding show commands and/or their outputs in support of this feature.

show user-plane-service statistics tethering-detection

The following fields are displayed in support of this feature:

- Current Tethered Subscribers
- Total Tethered Subscribers
- Total flows scanned

- Total Tethered flows detected
- Total Tethered flows recovered
- Total flows bypassed for scanning
- Tethering Detection Statistics (ip-ttl)
  - Total flows scanned
  - Tethered flows detected
  - Tethered uplink packets
  - Tethered downlink packets

show user-plane-service statistics rulebase name rulebase_name

The following fields are displayed in support of this feature:

- Tethering Detection (ip-ttl)
  - Total flows scanned
  - Tethered flows detected
  - Tethered uplink packets
  - Tethered downlink packets

RTP Dynamic Flow Detection

The **rtp dynamic-flow-detection** CLI command, under the ACS Rulebase Configuration mode, enables the Real Time Streaming Protocol (RTSP) and Session Description Protocol (SDP) analyzers to detect the child RTP and RTCP flows. If you configure the RTSP/SIP and SDP analyzers, and **rtp dynamic-flow-detection** CLI is present, then there's no need for configuring RTP/RTCP explicitly. With the **rtp dynamic-flow-detection** CLI command, the child RTP or RTCP flows get corelated to their parent RTSP/SIP-SDP flows.

Once the parent flow (RTSP/SIP-SDP) gets cleared, the child RTP/RTCP flows also gets cleared. In the absence of this CLI, the L7 layer analysis for RTP and RTCP needs a separate analyzer configuration. There’s no correlation of RTP/RTCP flows to RTSP/SIP-SDP flow.

Rule-matching for Bearer-specific Filters

The Rule-matching for Bearer-specific Filters functionality includes:

- IMSI-based rules are matched as per the subscribers IMSI.
- APN-based rules allow you to define rule expressions to match Access Point Name (APN) of the bearer flow.
- RAT-Type allows you to define rule expressions to match Radio Access Technology (RAT) in the bearer flow.

Configuring IMSI Pool

Use the following CLI commands to configure the IMSI pool.

```
configure
   active-charging service service_name
      imsi-pool pool_name
         imsi { imsi_number | range start_imsi to end_imsi }
```

The imsi-pool can contain either IMSI value or range of IMSI.

Configuring Rule-line ACS Ruledef

Use the following CLI commands to configure rule-line in ACS Ruledef Configuration mode.

```
configure
   active-charging service service_name
      ruledef ruledef_name
         bearer 3gpp imsi { = imsi_value } | { range imsi-pool pool_name }
         bearer 3gpp apn operator apn_name
         bearer 3gpp rat-type operator rat_type
```

IMSI range can be configured in a rule with the help of IMSI pool.

bearer 3gpp rat-type *operator rat_type*:

- *operator* must be one of the following:
  - **!=** : Does not equal
  - **=** : Equals
- **NOTE**: In this release, **wlan** is the qualified *rat_type*.

Configuring HTTP Content-Type

Use the following CLI commands to define rule expressions to match value in HTTP Content-Type entity-header field.

```
configure
   active-charging service service_name
      ruledef ruledef_name
         http content type [ case-sensitive ] operator content_type
```

Show CLIs

Use the following CLI on UPF to see information about IMSI pool that is configured in a service:**show user-plane-service imsi-pool name** *pool_name*

Enablement of IP Source Violation Detection for Uplink Packets

Table 41: Feature History

Feature Name	Release Information	Description
Enable IP Source Violation Detection for Uplink Packets	2024.03.0	This feature on cnSGWc lets UPF check if the origin of the uplink packets is from the right subscriber. This feature provides the following support: • Enhance the security and privacy of the subscribers by preventing the leakage of their data to unauthorized parties. This feature further reduces the risk of legal and regulatory issues for the service provider by complying with the lawful interception requirements. • Delayed reuse of the Tunnel Endpoint Identifier (TEID) that cnSGW assigns for a session until the session termination prevents tapping the wrong subscriber, risk of packet misassociation, and reporting by Lawful Interception (LI) agencies. This feature uses the existing CLI command **show subscribers user-plane-only full all | callid** *callid* to view the configured **IP Source Violation Action** option. **Default Setting**: Disabled – Configuration Required to Enable

The IP source violation detection feature is enabled on cnSGW. This information is communicated from cnSGWc to UPF. Based on the configuration done in cnSGW, which is to ignore or discard the packets, the User Plane acts on the source violated packet.

After a UE connected to a network, cnSGW assigns a unique Tunnel Endpoint Identifier (TEID) for a new session and includes TEID in the Create Session Request message to UPF. The UPF acknowledges the session establishment and uses the TEID to route user data packets.

With this feature, the Tunnel Endpoint Identifier (TEID) that the cnSGW assigns for a session per subscriber, is not reused immediately for another subscriber until the session termination. The delayed TEID reuse avoids tapping the wrong subscriber and reporting by LI agencies.

You can verify the IP source violation on User Plane through the **show subscribers user-plane-only full all | callid** *call_id* CLI command.

For details on the validating the uplink packets for IP source violation, see Validation of Uplink Packets for IP Source Violation in the *cnSGWc Configuration and Admininstration Guide*.

How Enabling of IP Source Violation for Uplink Packets Works

Following are the key points when you enable or disable this feature:

- By default, this feature is disabled.
- When this feature is enabled with the IGNORE option, the packets are validated for invalid source IP address. The corresponding statistics are incremented accordingly. However, there is no action on the packet because of the invalid IP.
- When this feature is enabled with the DISCARD option, the packets with invalid source IP address are discarded. The corresponding statistics are incremented accordingly.

Verifying IP Source Violation for Uplink Packets

Use this procedure to verify on UPF if the IP Source Violation feature is enabled or disabled on cnSGWc along with the configured values.

Procedure

Enter the **show subscribers user-plane-only full all | callid <>**.

Example:

```
show subscribers user-plane-only full all | callid <>
```

A new counter **ip source violations ignored:<>** is added to show the number of IP source violated packets that have been ignored and forwarded.

The following is an example output of the **show subscribers user-plane-only full all | callid <>** CLI command where the values for **IP Source Violation Action**, **ip source violations**, and **ip source violations ignored** are configured.

```
show subscribers user-plane-only full all
  Local SEID      : [0x0004000000000000] 1125899906842624
  Remote SEID     : [0x0004000000000005] 1125899906842629
  State           : Connected
  Connect Time    : Wed Jun  5 05:09:22 2024
  Idle time       : 00h11m19s
  Access Type: n/a                      Network Type: IP
  user-plane-service-name: UP21
  active-service-scheme-name:
  Callid: 00004e21
  Rulebase:
  Interface Type: Sxa
  eMPS Session: No
  eMPS Session Priority: 0
  Session-Type: normal
  Precedence-order: 0
  Data Queue: 0
  Card/Cpu: 1/0                         Sessmgr Instance: 1
  IP address: n/a
  Next Hop Ip Address: n/a
  Source context: EPC2-UP               Destination context: EPC2-UP
  PDN-Instance: n/a
  User-plane-Sx-addr: 20.20.20.106
  Control-plane-Sx-addr: 20.20.20.101
  Number of associated PDRs : 2
  Number of associated ADC PDRs : 0
  Number of associated FARs : 2
```

```
Number of associated QERs : 0
Number of associated BARs : 0
Number of associated URRs : 1
Uplink APN AMBR (bps) : 0       Downlink APN AMBR (bps) : 0
active input acl: n/a                active output acl: n/a
active input ipv6 acl: n/a             active output ipv6 acl: n/a
Bandwidth Policy: n/a
FW-and-NAT Policy: n/a
FW-and-NAT Policy ID: n/a
Firewall Policy IPv4: n/a
Firewall Policy IPv6: n/a
NAT Policy NAT44: n/a
NAT Policy NAT64: n/a
Converged Session: No                 Converged Peer Callid:      n/a
Visited Call: No
IP Source Violation Action: DISCARD
Subscriber Parameters:
IMSI: 404005123456789
IMEI: 112233445566778
MSISDN: 9890098900
Rat Type: 6
Old Rat Type: n/a
CF Policy ID: n/a
GX CF Policy ID: n/a
Roaming Status: n/a
Old Roaming Status: n/a
input pkts: 5                                 output pkts: 0
input bytes: 250                              output bytes: 0
input bytes dropped: 123                      output bytes dropped: 0
input pkts dropped: 3                         output pkts dropped: 0
CF Buffered Uplink Packets: 0                 CF Buffered Downlink Packets: 0
CF Buffered Uplink Bytes: 0                   CF Buffered Downlink Bytes: 0
Uplink Packets in Buffer: 0                   Downlink Packets in Buffer: 0
Buff Over-limit Uplink Pkts: 0                Buff Over-limit Downlink Pkts: 0
DDN buffered pkts : 0                         DDN buffered bytes : 0
DDN buffer overflow drop pkts : 0             DDN buffer overflow drop bytes : 0
pk rate from user(bps): 0                     pk rate to user(bps): 0
ave rate from user(bps): 0                    ave rate to user(bps): 0
sust rate from user(bps): 0                   sust rate to user(bps): 0
pk rate from user(pps): 0                     pk rate to user(pps): 0
ave rate from user(pps): 0                    ave rate to user(pps): 0
sust rate from user(pps): 0                   sust rate to user(pps): 0
ipv4 bad hdr: 0                               ipv4 ttl exceeded: 0
ipv4 fragments sent: 0                        ipv4 could not fragment: 0
ipv4 bad length trim: 0
ipv4 input mcast drop: 0                      ipv4 input bcast drop: 0
input pkts dropped (0 mbr): 0                 output pkts dropped (0 mbr): 0
ipv4 input acl drop: 0                        ipv4 output acl drop: 0
ipv6 input acl drop: 0                        ipv6 output acl drop: 0
ip source violations: 3                       ipv4 output no-flow drop: 0
ipv6 bad hdr: 0                               ipv6 bad length trim: 0
ip source violations ignored: 0
ipv4 icmp packets dropped: 0
APN AMBR Uplink Pkts Drop: 0                   APN AMBR Downlink Pkts Drop: 0
APN AMBR Uplink Bytes Drop: 0                  APN AMBR Downlink Bytes Drop: 0
APN AMBR Uplink Pkts IP pref lowered: 0        APN AMBR Downlink Pkts IP pref lowered: 0

APN AMBR Uplink Bytes IP pref lowered: 0       APN AMBR Downlink Bytes IP pref lowered: 0

ITC Uplink Pkts Drop: 0                        ITC Downlink Pkts Drop: 0
ITC Uplink Bytes Drop: 0                       ITC Downlink Bytes Drop: 0
ITC Uplink Pkts IP pref lowered: 0             ITC Downlink Pkts IP pref lowered: 0
ITC Uplink Bytes IP pref lowered: 0            ITC Downlink Bytes IP pref lowered: 0
ITC Terminated Flows:: 0                       ITC Redirected Flows: 0
```

```
  Flow Action Terminated Flows: 0                  Flow Action Redirected Flows: 0
  ToS marked Uplink Pkts: 0                        ToS marked Downlink Pkts: 0
  CC Dropped Uplink Pkts: 0                        CC Dropped Downlink Pkts: 0
  CC Dropped Uplink bytes: 0                       CC Dropped Downlink Bytes: 0
  Uplink Inflight Pkts: 0                          Downlink Inflight Pkts: 0
  QFI Mismatch Uplink Pkts Drop: 0                 QFI Mismatch Downlink Pkts Drop: 0
  Total QoS-Group(s) Active: 0
  DNS-to-EDNS Uplink Pkts: 0                     DNS-to-EDNS Uplink Bytes: 0
  EDNS Response Received: 0

  Flow information:
    Current Active Flows:
      TCP: 0
      UDP: 0
    Total Flows:
      TCP: 0
      UDP: 0
      FP:  0

Static & Predef Rule Match stats:
Rule Name             Pkts-Down  Bytes-Down Pkts-Up    Bytes-Up   Hits       Match-Bypassed
FP-Down(Pkts/Bytes)  FP-Up(Pkts/Bytes)
--------------------- ---------- ---------- ---------- ---------- ---------- --------------
-------------------  ------------------

Dynamic Rule Match stats:
PDR Id   Rule Name            Pkts-Down  Bytes-Down Pkts-Up    Bytes-Up   Hits       Match-Bypassed
  FP-Down(Pkts/Bytes)  FP-Up(Pkts/Bytes)
-------- -------------------- ---------- ---------- ---------- ---------- ---------- --------------
  -------------------  ------------------
0x0002   N/A                         0          0          5        250          0              0
      0/0               5/250

Post-Processing Rule Match stats:
Rule Name             Pkts-Down  Bytes-Down Pkts-Up    Bytes-Up   Hits       Match-Bypassed
FP-Down(Pkts/Bytes)  FP-Up(Pkts/Bytes)
--------------------- ---------- ---------- ---------- ---------- ---------- --------------
-------------------  ------------------

Firewall-NAT Access Rule Match stats:
Firewall-Ruledef Name Pkts-Down Bytes-Down    Pkts-Up   Bytes-Up Dr-Pkts-Dn Dr-Byts-Dn Dr-Pkts-Up
Dr-Byts-Up       Hits
--------------------- ---------- ---------- ---------- ---------- ---------- ---------- ----------
---------- ----------

QoS-Group Statistics:
QGR Name              Pkts-Down  Bytes-Down Pkts-Up    Bytes-Up   Hits       Match-Bypassed
FP-Down(Pkts/Bytes)  FP-Up(Pkts/Bytes)
--------------------- ---------- ---------- ---------- ---------- ---------- --------------
-------------------  ------------------

SGACL Match stats:
ACL Name              Pkts-Down  Bytes-Down Pkts-Up    Bytes-Up   Pkts dropped
--------------------- ---------- ---------- ---------- ---------- -------------

  CIoT APN Rate Control:
    Allowed UL Limit:   0                        Allowed DL limit:    0
    Remaining UL Limit: 0                        Remaining DL limit:  0
    Allowed Time unit: unrestricted              Status Validity Time: N/A

Total subscribers matching specified criteria: 1
```

URL Blockedlisting

Feature Description

The URL blockedlisting feature regulates the subscriber's access to view or download content from websites whose URL or URI has been blockedlisted. It uses a database that records a list of URLs that indicates if the detected URL is categorized to be blocked or not.

How it Works

To enable the URL blockedlisting feature on UPF, URL blockedlisting database should be present with a name "optblk.bin" under flash, or SFTP or under its subdirectory. This database directory path must be configured on user-plane, after user-plane services are brought up.

HTTP Analyzer must be enabled for URL blockedlisting. The HTTP analyzer extracts URL information from the incoming HTTP request data packet. Extracted URL content is compared with the URL Blockedlisting database. When the URL of incoming HTTP data packet matches with the database URL entry, that URL is treated as blockedlisted URL and one of the following actions takes place on that HTTP packet:

- Termination of flow
- Packet is discarded

The URL blockedlisting configurations must be configured under Rulebase configuration in Active Charging Service. Also, two URL blockedlisting methods – Exact and Generic, are supported at Active Charging Service-level configuration.



Important Blockedlisting database(s) are provided by – Internet Watch Foundation (IWF) and National Center for Missing and Exploited Children (NCMEC). The UPF always receives the blockedlisting database in Optimized Format.

URL Blockedlisting Database Upgrade

URL database upgrade is supported in following two ways:

- Timer-based upgrade or Auto upgrade
- CLI-based upgrade or Manual upgrade

Timer-based or Auto-upgrade

After the database is loaded on the chassis for the first time, a timer, for a duration of 5 minutes, is started. This process is started to auto upgrade the database.

If at the expiry of the timer, a valid database with higher version is available at the directory path, then database upgrade procedure is initiated, and a newer version of the database is loaded on the UPF.

To upgrade a URL blockedlisting database, a higher version of valid URL Blockedlisting database with name "optblk_f.bin" should be present at same directory as that of current database "optblk.bin".

After the database is upgraded successfully, the earlier "optblk.bin" file gets renamed as "optblk_0.bin" and "optblk_f.bin" file gets renamed as "optblk.bin". Here, "optblk_0.bin" file is treated as a backup file of older database.

If an additional upgrade is performed, then “optblk_0.bin” file will be renamed as “optblk_1.bin” file and current “optblk.bin” will get renamed as “optblk_0.bin”, and so on.

See the *Loading URL Blockedlisting Database on UPF* section to configure the number of backup files to be stored in the database.

CLI-based or Manual Upgrade

See the *Upgrading the URL Blockedlisting Database* section to upgrade the current database to a newer version.

Configuring URL Blockedlisting

Loading URL Blockedlisting Database on UPF

Use the following configuration to load URL blockedlisting database on UPF.

In releases prior to 2022.01.0:

```
configure
  url-blacklisting database directory path database_directory_path
  url-blacklisting database max-versions max_version_value
  end
```

From 2022.01.0 and later releases:

```
configure
  url-blockedlisting database directory path database_directory_path
  url-blockedlisting database max-versions max_version_value
  end
```

NOTES:

- **database directory path**: Configures the database directory path.

  The *database_directory_path* is a string of size 1 to 255.

- **max-versions**: Configures the maximum database upgrade versions.

  The *max_version_value* is an integer from 0 to 3.

Upgrading the URL Blockedlisting Database

Use this configuration to manually upgrade the URL blockedlisting database.

In releases prior to 2022.01.0:

```
upgrade url-blacklisting database
end
```

From 2022.01.0 and later releases:

```
upgrade url-blockedlisting database
end
```

Configuration to Enable URL Blockedlisting

Use the following configuration to enable URL blockedlisting feature on UPF.

In releases prior to 2022.01.0:

```
configure
  active-charging service service_name
    url-blacklisting match-method [ exact | generic ]
    rulebase rulebase_name
      url-blacklisting action [ discard | terminate-flow ]
      end
```

From 2022.01.0 and later releases:

```
configure
  active-charging service service_name
    url-blockedlisting match-method [ exact | generic ]
    rulebase rulebase_name
      url-blockedlisting action [ discard | terminate-flow ]
      end
```

NOTES:

- **match-method [exact | generic]**: Specifies the match method used for URL blockedlisting.
  - **exact**: URL Blockedlisting perform an exact-match of URL.
  - **generic**: URL Blockedlisting perform generic-match of URL.
- **url-blockedlisting action [discard | terminate-flow]**:
  - **discard**: Discards the HTTP packet received.
  - **terminate-flow**: Terminates the flow of the HTTP packet received.

Monitoring and Troubleshooting

This section provides information regarding the CLI command available in support of monitoring and troubleshooting the feature.

Show Command(s) and/or Outputs

This section provides information regarding show commands and/or their outputs in support of this feature.

show user-plane-service url-blacklisting database

The following fields are displayed in support of this feature:

- URL Blacklisting Static Rating Databases:
  - Last Upgrade Status
  - Path
    - Database Status
    - Number of URLs in DB
    - Type
    - Version
    - Creation Time

- Hostname
- Comment
- Last Access Time
- Last Modification Time
- Last Status Change Time

show user-plane-service url-blacklisting database url database_directory_path

The following fields are displayed in support of this feature:

- URL Blacklisting Static Rating Databases:
  - Last Upgrade Status
  - Path
    - Database Status
    - Number of URLs in DB
    - Type
    - Version
    - Creation Time
    - Hostname
    - Comment
    - Last Access Time
    - Last Modification Time
    - Last Status Change Time

show user-plane-service url-blacklisting database facility sessmgr all

The following fields are displayed in support of this feature:

- URL-Blacklisting SessMgr Instance Based Database Configuration
  - SessMgr Instance
  - BL DB Load Status
  - BL DB Version
  - Number of URLs
  - Checksum

show user-plane-service rulebase name rulebase_name

The following fields are displayed in support of this feature:

- URL-Blacklisting Action
- URL-Blacklisting Content ID

show user-plane-service inline-services info

The following fields are displayed in support of this feature:

- URL-Blacklisting: Enabled
  - URL-Blacklisting Match-method: Generic

show user-plane-service inline-services url-blockedlisting statistics

The following are displayed in support of this feature:

- Cumulative URL-Blockedlisting Statistics
  - Blockedlisted URL hits
  - Blockedlisted URL misses
  - Total rulebases matched

show user-plane-service inline-services url-blacklisting statistics rulebase name rulebase_name

The following fields are displayed in support of this feature:

- Rulebase Name
  - URL-Blacklisting Statistics
  - Blacklisted URL hits
  - Blacklisted URL misses
- Total rulebases matched

Configuring the Static and Pre-Defined Rules

This section describes how to configure the static and pre-defined rules under the charging action configuration.

```
configure
   active-charging service service_name
      charging-action charging_action_name
         flow action { discard [ downlink | uplink ] | redirect-url
redirect_url | terminate-flow }
         end
```

NOTES:

- **flow action**: Specifies the action to take on packets that match rule definitions.
  - **discard [downlink | uplink]**: Specifies to discard downlink or uplink packets.
  - **redirect-url** *redirect_url*: Specifies the URL to be redirected. For example, http://search.com/subtarg=#HTTP.URL#
  - **terminate-flow**: Specifies to terminate the flow.
- For redirect-url, configure HTTP analyzer under rulebase. Example:

```
route priority 70 ruledef http-port analyzer HTTP
   ruledef http-port
   tcp either-port = 80
   rule-application routing
   exit
```

Configuring ACS Ruledef for L7 Protocols for DPI

A ruledef represents a set of matching conditions across multiple L3 – L7 protocol based on protocol fields and state information. Each ruledef can be used across multiple rulebases within the active charging service.

Note In UPF, if rule-line addition or deletion inside a ruledef is done during active calls and data flows, then this configuration change is not applied for current flows. However, the configuration change applies to new calls and new flows on same calls.

The following is a sample configuration that describes how to create, configure, or delete ACS rule definitions.

```
configure
   active-charging service service_name
      ruledef ruledef_name
         dns { any-match value | query-type query_type | query-name query_name
 }
         ip any-match [ = | != ] [ TRUE | FALSE ]
         ip dst-address { operator { { ipv4_address | ipv6_address } | {
ipv4_address/mask | ipv6_address/mask} | address-group ipv6_address } | { !range |
 range } host-pool host_pool_name }
         ip server-ip-address { operator { { ipv4_address | ipv6_address } | {
ipv4_address/mask | ipv6_address/mask} | address-group ipv6_address } | { !range |
 range } host-pool host_pool_name }
         multi-line-or all-lines
         rule-application { charging | post-processing | routing }
         { tcp | udp } { either-port port_number }
         end
```

NOTES:

- **ruledef** *ruledef_name*: Specifies the ruledef to add, configure, or delete. *ruledef_name* must be the name of an ACS ruledef, and must be an alphanumeric string of 1 to 63 characters, and can contain punctuation

characters. Each ruledef must have a unique name. Host pool, port map, IMSI pool, and firewall, routing, and charging ruledefs must have unique names.

- If the named ruledef does not exist, it is created, and the CLI mode changes to the ACS Ruledef Configuration Mode wherein the ruledef can be configured.
- If the named ruledef already exists, the CLI mode changes to the ACS Ruledef Configuration Mode for that ruledef. The ACS Ruledef Configuration Mode is used to create and manage rule expressions in individual rule definitions (ruledefs).
- **ip any-match [= | !=] [TRUE | FALSE]:** This command defines the rule expressions to match IPv4/IPv6 packets. The *operator* and *condition* in the command specifies the following:
  - *operator*
    - !=: Does not equal
    - < =: Equals
  - *condition*
    - FALSE
    - TRUE
- **ip dst-address {** *operator* **{ {** *ipv4_address* **|** *ipv6_address* **} | {** *ipv4_address/mask* **|***ipv6_address/mask* **} | address-group** *ipv6_address* **} | { !range | range } host-pool** *host_pool_name* **}**: This command allows defining rule expressions to match IP destination address field within IP headers.
  - *ipv4_address* **|** *ipv6_address*: Specifies the IP address of the destination node for outgoing traffic. *ipv4_address* **|** *ipv6_address* must be an IP address in IPv4 dotted-decimal or IPv6 colon-separated-hexadecimal notation.
  - *ipv4_address/mask* **|** *ipv6_address/mask*: Specifies the IP address of the destination node for outgoing traffic. *ipv4_address/mask* **|** *ipv6_address/mask* must be an IP address in IPv4 dotted-decimal or IPv6 colon-separated-hexadecimal notation with subnet mask bit. The mask bit is a numeric value which corresponds to the number of bits in the subnet mask.
  - *address-group ipv6_address*: Specifies a group of IPv6 addresses configured with wildcard input and/or specialized range input. Multiple wildcard characters can be accepted as input and only one 2 byte range input will be accepted. Both wildcard character input and 2-byte range input can be configured together within a given IPv6 address.
  - **host-pool** *host_pool_name*: Specifies the name of the host pool. *host_pool_name* must be an alphanumeric string of 1 to 63 characters.
  - The *operator* in the command specifies the following:
    - !=: Does not equal
    - <: Lesser than or equals
    - =: Equals
    - >=: Greater than or equals

- **multi-line-or all-lines**: This command allows a single ruledef to specify multiple URL expressions. When a ruledef is evaluated, if the multi-line-or all-lines command is configured, the logical OR operator is applied to all the rule expressions in the ruledef to decide if the ruledef matches or not. If the multi-line-or all-lines command is not configured, the logical AND operator is applied to all the rule expressions.
- **rule-application { charging | post-processing | routing }**: This command specifies the rule application for a rule definition**.**
  - **charging**: Specifies that the current ruledef is for charging purposes.
  - **post-processing**: Specifies that the current ruledef is for post-processing purposes. This enables processing of packets even if the rule matching for them has been disabled.
  - **routing**: Specifies that the current ruledef is for routing purposes. Up to 256 rule definitions can be defined for routing in an Active Charging Service. Default: Disabled.
- **dns { any-match** *value* **| query-type** *query_type* **| query-name** *query_name* **}**: This command allows you to define rule expressions to match all DNS packets, or packets based on the query type or query name.

  ip server-ip-address *ip_address_value*: This command allows you to define rule expressions to match the IP address of the destination end of the connection.
- **{ tcp | udp } { either-port** *port_number* **}**: This command allows you to define rule expressions to match either a destination or source port number in UDP/TCP headers.

Charging Action Configuration for L7 Protocols for DPI

This section describes how to configure charging action. The charging action represents actions to be taken when a configured rule is matched. Actions could range from generating an accounting record (for example, an EDR) to dropping the IP packet, and so on. The charging action will also determine the metering principle—whether to count retransmitted packets and which protocol field to use for billing (L3, L4, L7, and so on).

The charging action configuration is used to define the QoS and charging related parameters associated with ruledefs.

```
configure
   active-charging service service_name
      charging-action charging_action
       allocation-retention-priority priority  [ pci pci_value  | pvi pvi_value

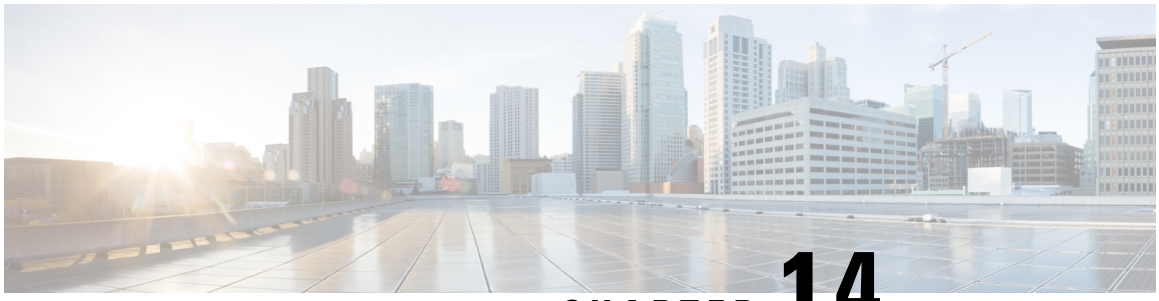
        billing-action egcdr
       cca charging credit [ rating-group coupon_id ] [ preemptively-request
 ]
        content-id content_id
        flow action { discard [ downlink | uplink ] | redirect-url
redirect_url | terminate-flow }
        flow limit-for-bandwidth { { direction { downlink | uplink }
peak-data-rate bps peak-burst-size bytes violate-action { discard |
lower-ip-precedence } [ committed-data-rate bps committed-burst-size bytes
 [ exceed-action { discard | lower-ip-precedence } ] ] } | { id id } }
```

```
        nexthop-forwarding-address ipv4_address/ipv6_address
        qos-class-identifier qos_class_identifier
        service-identifier service_id
        tft packet-filter packet_filter_name
        tft-notify-ue
        tos { af11 | af12 | af13 | af21 | af22 | af23 | af31 | af32 |
af33 | af41 | af42 | af43 | be | ef | lower-bits tos_value } [ downlink |
uplink ]
```

NOTES:

- **charging-action** *charging_action_name*: Specifies the name of a charging action. *charging_action_name* must be an alphanumeric string of 1 to 63 characters and can contain punctuation characters. Each charging action must have a unique name.
- If the named charging action does not exist, it is created, and the CLI mode changes to the ACS Charging Action Configuration Mode wherein the charging action can be configured.
- If the named charging action already exists, the CLI mode changes to the ACS Charging Action Configuration Mode for that charging action.
- **allocation-retention-priority** *priority* **[pci** *pci_value* **| pvi** *pvi_value* **]**: Configures the Allocation Retention Priority (ARP). *priority* must be an integer value in the range of 1-15.
  - **pci** *pci_value* : **Specifies the Preemption Capability Indication (PCI) value. The options are:**
    - MAY_PREEMPT - Flow can be preempted. This is the default value.
    - NOT_PREEMPT - Flow cannot be preempted
  - **pvi** *pvi_value*: Specifies the Preemption Vulnerability Indication (PVI) value. The options are:
    - NOT_PREEMPTABLE - Flow cannot be preempted. This is the default value.
    - PREEMPTABLE - Flow can be preempted
- **billing-action**: Configures the billing action for packets that match specific rule definitions.
- **cca charging credit**: Enables or disables credit control charging credit behaviour.
- **content-id**: Configures the rating group.
- **flow action**: Specifies the action to take on packets that match rule definitions.
  - **discard [downlink | uplink]**: Specifies to discard downlink or uplink packets.
  - **redirect-url** *redirect_url*: Specifies the URL to be redirected.
  - **terminate-flow**: Specifies to terminate the flow.
- **flow limit-for-bandwidth**: Configures the QoS parameters such as MBR, GBR, and so on.
  - peakdatarate(MBR): Default is 3000 bps
  - peakburstsize: Default is 3000 bytes
  - committedDataRate(GBR): Default is 144000 bps
  - committedBurstSize: Default is 3000 bytes

- **nexthop-forwarding-address** *ipv4_address/ipv6_address* ,: Configures the nexthop forwarding address.
- **qos-class-identifier** *qos_class_identifier* : Configures the QCI for a charging action. *qos_class_identifier* must be an integer value in the range of 1-9 or from 128-254 (Operator specific).
- **service_identifier** *service_id*: Configures the service identifier to use in generated billing records.*service_id* must be an integer value in the range of 1-2147483647.
- **tft packet-filter** *packet_filter_name*: Specifies the packet filter to add or remove from the current charging action. *packet_filter_name* must be the name of a packet filter, and must be an alphanumeric string of 1 to 63 characters.
- **tft-notify-ue**: Control the TFT updates towards the UE based on certain trigger conditions.
- **tos**: Configures the Type of Service (ToS) octets.



CHAPTER 14

Device ID in EDNS0 Records

- Feature Summary and Revision History, on page 189
- Feature Description, on page 190
- How it Works, on page 190
- Configuring EDNS Format and Trigger Action, on page 193
- Monitoring and Troubleshooting, on page 195

Feature Summary and Revision History

Summary Data

Table 42: Summary Data

Applicable Product (s) or Functional Area	5G-UPF
Applicable Platforms	VPC-SI SMI
Feature Default Setting	Disabled – Configuration Required
Related Changes in this Release	Not Applicable
Related Documentation	*UCC 5G UPF Configuration and Administration Guide*

Revision History

Table 43: Revision History

Revision Details	Release
TCP support is added.	2021.02.1
First introduced.	2021.01.2

Feature Description

The Device ID in EDNS0 offers each enterprise with a customized domain blocking through Umbrella. To enable this functionality:

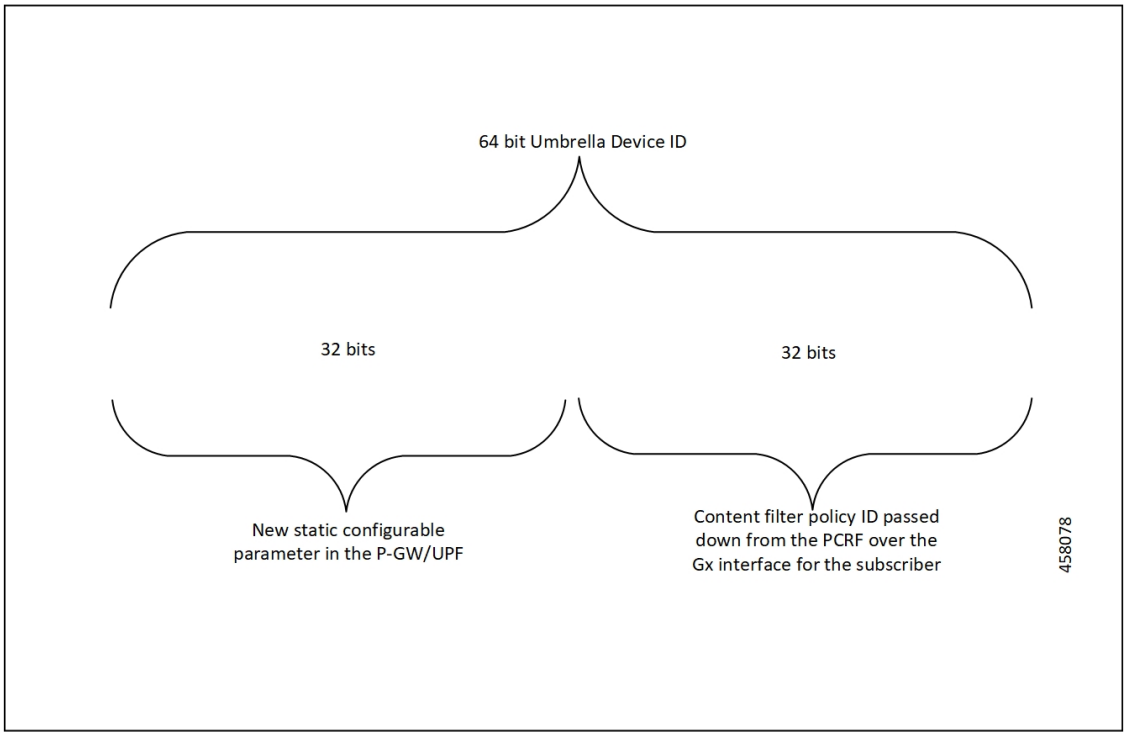
- The UPF must reformat a subscriber DNS request into an EDNS0 request, and.
- The UPF must include an Umbrella “Device ID” in the EDNS0 packet so that the Umbrella DNS resolver can use the Device ID to apply the domain filter associated or configured with the Device ID in the EDNS0 packet.

Presently, the Session Management Function (SMF) receives the domain filtering policy ID from PCRF/PCF. The SMF passes the domain filtering policy ID to the User Plane Function (UPF) in the Subscriber Parameters. The UPF uses the domain filtering policy ID to apply domain filtering functionality to the subscriber.

How it Works

New CLIs are introduced to configure and trigger the EDNS0 functionality.

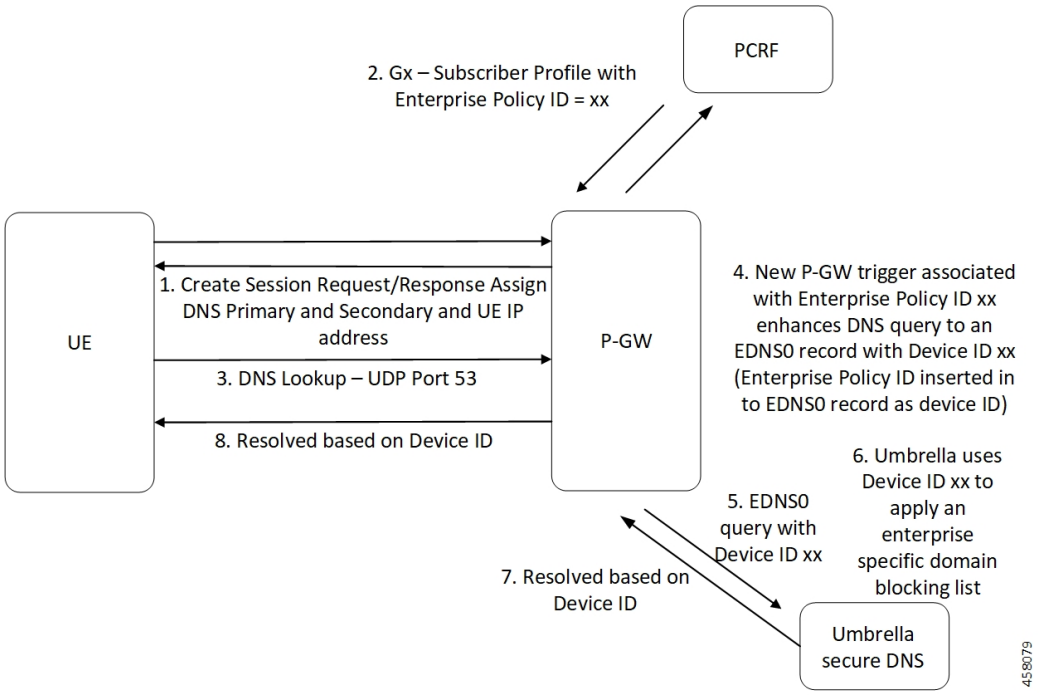
The EDNS0 packet receives the 64-bit device ID as OPT RR data. The first 32 bits of all device IDs is a fixed value that is configured in the UPF. The last 32 bits of a subscriber device ID is the content filter ID value received from the PCRF/PCF. The UPF concatenates the two 32-bit values to build a subscriber full 64-bit Device ID for populating in the subscriber EDNS0 queries. New CLI helps to configure the first 32 bit of static device-id value. If you don’t configure the 32-bit static prefix CLI, the outgoing packet shows the device-id = 32-bit CF PolicyID.



The Device ID number in the EDNS0 record allows the Umbrella DNS system to apply a custom set of domain filters for the EDNS0 queries.

Process Flow

The following process flow describes about the Content Filtering enhancement to insert Device ID in EDNS0 records:



EDNS0 Packet Format

The enterprise policy ID (CF_POLICY_ID) from PCRF helps to create the Device ID. The SMF sends the device ID to the UPF. Adding the Device ID to the DNS packet helps in creating the EDNS0 packet. The format of EDNS0 packets is specified by RFC2671. The following are few specifics:

- Following is the structure for the fixed part of an OPT RR:

```
Field Name   Field Type     Description
------------------------------------------------------
NAME         domain name    empty (root domain)
TYPE           u_int16_t          OPT
CLASS          u_int16_t          sender's UDP payload size
TTL            u_int32_t          extended RCODE and flags
RDLEN          u_int16_t          describes RDATA
RDATA          octet stream       {attribute, value} pairs
```

- Following is the variable part of an OPT RR encoded in its RDATA:

```
     +0 (MSB)                            +1 (LSB)
    +---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
 0: |                         OPTION-CODE                           |
    +---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
 2: |                         OPTION-LENGTH                         |
```

```
   +---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
4: |                                                               |
   /                       OPTION-DATA                             /
   /                                                               /
   +---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
```

- OPTION-CODE: Assigned by IANA
- OPTION-LENGTH: Size (in octets) of OPTION-DATA
- OPTION-DATA- Varies per OPTION-CODE

Example: If received policy-id from PCF/PCRF is “1234” and static prefix configured on UPF is “5678”.

64-bits Device-ID will be “0000162e000004d2”.

- 0000162e -- 5678 (Decimal)
- 000004d2 -- 1234 (Decimal)

RDATA 69 42 00 0f 4f 70 65 6e 44 4e 53 00 00 16 2e 00 00 04 d2

- 6942 -- option-code
- 000f -- option-length
- 4f70656e444e53 -- OpenDNS (String)
- 0000162e -- 5678 (MSB)
- 000004d2 -- 1234 (LSB)

EDNS0 with IP Readdressing

The new CLI is configured within trigger action to readdress the DNS traffic to the Umbrella DNS. This CLI uses the existing readdress server list configuration from the ACS service. Readdressing of packets based on the destination IP address of the packets enables redirecting gateway traffic to configured server/port in the readdressed server list.

Behavior and Restrictions

Following are the behavior and restrictions applicable for this feature:

- Trigger Condition is evaluated at flow creation time. Any change in trigger condition in between the flow doesn’t affect the existing flow but affects the new flows.
- Any change to trigger action is applicable on the same flow.
- Neither CF nor EDNS is enforced when the CF Policy ID range is defined but Service-schema is not defined, or the Trigger condition pertaining to EDNS is not configured.
- If no CF Policy ID is received from Gx, range check is not performed, and content filtering works as defined in rule base.
- Cases where the ‘security-profile’ CLI is not associated with the ‘EDNS format’ CLI in Trigger Action, the device-id in the outgoing EDNS packet is sent with only 32-bit CF Policy ID.

- DNS queries with type other than A, AAAA, CNAME, NS, PTR, SRV, TXT, NULL are not to be EDNS converted.
- CF Policy ID change over Gx in between inflow are not applicable for the current flows. The current flows continue to insert the CF Policy ID present at the time of flow creation.

Limitation

Following are the limitations for this feature:

- The feature doesn't support the EDNS response packet reformat.
- The UPF must be able to include the IMSI MSISDN tag value in the EDNS0 queries. This feature doesn't support the encrypted IMSI in EDNS0 packet. This feature also doesn't support the following configuration on the EDNS fields currently.

```
configure
   active-charging-service  service_name
      edns
         fields  fields_name
            tag default device-id
            tag 101 imsi encrypt
            tag 102 pgw-address
            end
```

Configuring EDNS Format and Trigger Action

Use the following configuration to enable DNS filtering:

```
configure
   active-charging-service service_name
      content-filtering range start_min_val to end_max_val
```

If the range parameter is set from 10 through 1000, any subscriber profile with a content filter policy ID from 10 through 1000 uses the standard content filtering functionality. Any subscriber profile with a content filter policy ID higher than 1000 or lower than 10 triggers the new EDNS0 functionality.

Use the following configuration to disable DNS filtering:

```
configure
   active-charging-service service_name
      no content-filtering range
```

When DNS filtering is disabled, the standard content filtering policies resume as configured or as received from the PCF.

Use the following configuration to configure the EDNS packet action and format under the active-charging service:

```
configure
   active-charging-service service_name
      trigger-condition trigger_condition_name
         external-content-filtering
```

```
      app-proto = dns
    end
```

NOTES:

- **external-content-filtering**: Enables EDNS0 feature. When this flag is true along with the range criteria, EDNS0 feature is enabled. By default, this flag is disabled.
- **app-proto** = **dns**: Avoids the IP readdressing of the non-DNS traffic. If this CLI is enabled with multiline-or cli, then all DNS traffic is EDNS encoded.

The following configuration leads the trigger action to define the EDNS format to be inserted in the EDNS packet:

```
configure
  active-charging-service service_name
    trigger-action trigger_action_name
      edns-format format_name
        security-profile profile_name
          flow action readdress server-list server_list_name [ hierarchy ]
 [ round-robin ] [ discard-on-failure ]
           end
```

NOTES:

- **trigger-action** *trigger_action_name*: Enables you to configure the flow action CLIs in the trigger action.
- **edns-format** *format_name*: Use the EDNS format when EDNS is applied.
- **security-profile** *profile_name*: Defines the security profile configuration in the EDNS to add mapping with the Device-ID.

Note Device ID in EDNS0 Records feature supports multiple security profiles.

- **flow action readdress server-list** *server_list_name* **[hierarchy] [round-robin] [discard-on-failure]**: Associates the EDNS with IP readdressing. Use IP readdressing to readdress the packets to the configured server IPs. This CLI in trigger action supports only server list configuration. It doesn't support single-server IP or port configuration like charging action.

Use the following configuration to insert the CF policy ID in the EDNS:

```
configure
  active-charging-service service_name
    edns
      fields fields_name
        tag { val { imsi | msisdn | cf-policy-id }}
        end
```

To configure the 32 MS bit, static value is provided at the EDNS level with the security profile.

```
security-profile security_profile cf-policy-id-static-prefix value
```

Use the following configuration to insert a new tag specifying the payload length:

```
tag default payload-length [ tcp | udp ] value
```

Value range: 576–4096

Sample Configuration

Following is the sample configuration for configuring the EDNS packets:

```
configure
active-charging service ACS
content-filtering range 10 to 100
ruledef dns-port
udp either-port = 53
tcp either-port = 53
multi-line-or all-lines
rule-application routing
#exit
 readdress-server-list re_adr_list_ta
  server 209.165.202.141
  server 2001::14
  server 209.165.202.142
  server 2001::15
#exit
rulebase starent
route priority 20 ruledef dns-port analyzer dns
#exit
edns
security-profile sec_profile cf-policy-id-static-prefix 123456
fields test_fields
tag 26946 cf-policy-id
#exit
format test_format
fields test_fields encode
#exit
#exit
trigger-action TA1
edns format test_format security-profile sec_profile
flow action readdress server-list re_adr_list_ta hierarchy
#exit
trigger-condition TC1
external-content-filtering
app-proto = dns
#exit
service-scheme SS1
trigger flow-create
priority 1 trigger-condition TC1 trigger-action TA1
#exit
subs-class SC1
rulebase = starent
multi-line-or all-lines
#exit
subscriber-base SB1
priority 1 subs-class SC1 bind service-scheme SS1
exit
end
```

Monitoring and Troubleshooting

Following are the show commands and outputs in support of enhance content filtering support to Insert device ID in EDNS0 records.

Show Commands and Outputs

Following are the show commands and outputs that are modified in support of the enhance content filtering support to Insert device ID in EDNS0 records.

- **show user-plane-service inline-services info**

```
CF Range: Enabled <<<<
  Start Value:  1 <<<<
  End Value:    1000 <<<
```

- **show user-plane-service statistics analyzer name dns**: output is modified to include the "EDNS Response Received" in both "EDNS Over UDP" and "EDNS Over TCP" sections.
- **show subscribers user-plane-only full callid**: output is modified to include the following parameters in the EDNS statistics per subscriber.
  - DNS-to-EDNS Uplink Pkts
  - DNS-to-EDNS Uplink Bytes
  - EDNS Response Received
- **show user-plane-service edns all**

```
Fields:
    Fields Name: fields_1
    tag 26946 cf-policy-id

    Fields Name: fields_2
    tag 2001 imsi
    tag 2002 msisdn
    tag 26946 cf-policy-id

    Format:
    Format Name: format_1
    fields fields_1 encode

    Format Name: format_2
    fields fields_2 encode

    Security-profile Name: high
    CF Prefix Policy ID: 1234
```

Use the following show commands to view the Trigger Action statistics:

- **show user-plane-service statistics trigger-action all**
- **show user-plane-service statistics trigger-action name** *trigger_action_name*
- **show user-plane-service trigger-condition all**

```
Trigger-Condition: TC1
       External-content-filtering : Enabled
       App-proto : dns
       Multi-line-OR All lines : Disabled
```

- **show user-plane-service trigger-action all**

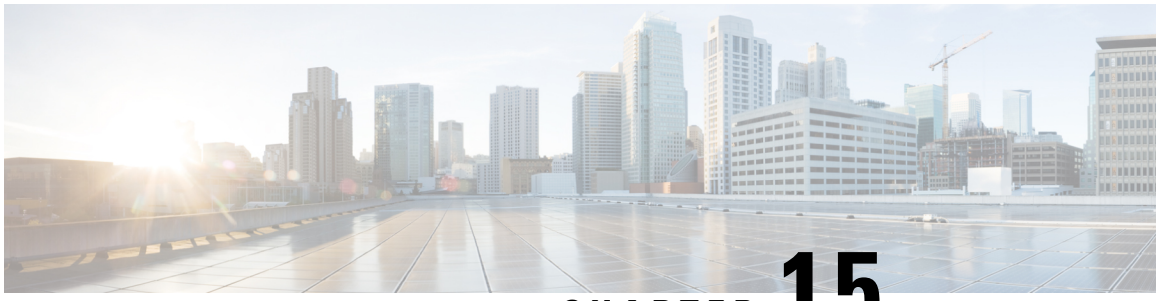
```
Trigger-Action: TA1
       HTTP Response Based TRM      : none
       HTTP Response Based Charging : none
       Throttle Suppress            : Disabled
```

```
Flow Recovery                : Disabled
Traffic Optimization         : Disabled
Step Up GBR                  : Disabled
Step Down GBR                : Disabled
TCP Acceleration             : Disabled
TCP Acceleration Threshold   : Disabled
Service-Chain                : none
UP-Service-Chain             : none
EDNS-Encode                  : Enabled
Flow-IP-Readdressing         : Enabled
```

Bulk Statistics

The following bulk statistics are available in support of the Device ID in EDNS0 Records feature:

SCHEMA: ECS	
Statistics	**Description**
ecs-dns-udp-edns-encode-succeed	The count of DNS to EDNS converted packets over UDP
ecs-dns-udp-edns-encode-failed	The count of failed DNS to EDNS conversions over UDP
ecs-dns-udp-edns-encode-response	The count of responses received for EDNS query over UDP
ecs-dns-tcp-edns-encode-succeed	The count of DNS to EDNS converted packets over TCP
ecs-dns-tcp-edns-encode-failed	The count of failed DNS to EDNS conversions over TCP
ecs-dns-tcp-edns-encode-response	The count of responses received for EDNS query over TCP



CHAPTER 15

Downlink Data Notification

- Feature Summary and Revision History, on page 199
- Feature Description, on page 200
- How It Works, on page 200
- DDN Throttling for non-Release 10 Compliant MME, on page 209
- DDN Throttling for Release 10 Compliant MME, on page 211
- Idle Timer for SAE-GW Sessions, on page 212
- S-GW Session Idle Timeout, on page 213
- Show Commands Input and/or Outputs, on page 213

Feature Summary and Revision History

Summary Data

Applicable Product(s) or Functional Area	5G-UPF
Applicable Platform(s)	VPC-SI SMI
Feature Default Setting	Disabled – Configuration Required
Related Changes in this Release	Not Applicable
Related Documentation	*UCC 5G UPF Configuration and Administration Guide*

Revision History

Revision Details	Release
First introduced.	2021.02.0

Feature Description

The Downlink Data Notification (DDN) messages with support for DDN Delay and DDN Throttling, and buffering in SAEGW when UE is in Idle State, is supported in 5G-UPF.

How It Works

This section provides an overview of how this feature works.

- Buffering is supported at SAEGW-U.
- Support of buffering starts when UE moves to IDLE state due to Release Access Bearer.
- ACTIVE to IDLE transition:
  - When the UE moves to ECM-IDLE state, since the SAEGW supports buffering capability and decides to activate buffering in SAEGW-U for the session, the SAEGW-C informs the SAEGW-U through an Sx session modification.
  - After the buffering starts, when the first downlink packet arrives on any bearer, the SAEGW-U informs the SAEGW-C. The SAEGW-U sends an Sx reporting message to the SAEGW-C, unless specified otherwise, and identifies the S5/S8 bearer on which the downlink packet is received.
  - On receiving the reporting message, the SAEGW-C decides whether to send a DDN message to the MME, as defined in 3GPP TS 23.401 [2]. The DDN notification is sent with the Sx-Usage-Report.
- IDLE to ACTIVE transition:
  - At the UE transition to ECM-CONNECTED state, the SAEGW-C updates the SAEGW-U through Sxa interface with the F-TEIDu of the eNodeB/RNC/SGSN. The buffered data packets, if any, are then forwarded to the eNodeB/RNC/SGSN by the SAEGW-U.
- If the Apply Action is BUFFER, and SGW-U recovers, the SGW-U initiates Sx Report (with DLDR Report Type) on arrival of the downlink data packet.
- In SGW-U, a timer is implemented that starts after each Sx Report (with DLDR report Type) is sent. If the Apply Action is not changed, then on timer expiry, Sx Report (with DLDR Report Type) gets initiated again.
- ARP of the bearer is included in the DDN message.
- In a multi-PDN session, if the DDN is initiated for one PDN and then data is received on another PDN, wherein the bearer has higher priority, then the DDN is initiated again with the higher priority ARP value.

Downlink Data Notification – Delay (DDN-D) Support

Under certain conditions, when UE triggers a service request, uplink and downlink data is triggered and is received at the SGW-C even before the Modify Bearer Request (MBR) is received causing unnecessary Downlink Packet Notification messages sent that increases the load in MME.

In such cases, the MME monitors the rate at which these events occur. If the rate becomes significant (as configured by the operator) and the MME's load exceeds an operator-configured value, the MME indicates "Delay Downlink Packet Notification Request" with parameter D to the Serving Gateway, where D is the requested delay given as an integer with multiples of 50 milliseconds, or zero. The S-GW then uses this delay in between receiving downlink data and sending the Downlink Data Notification message.

The Downlink Data Notifications are supported for both Collapsed and Pure-S calls.

Due to the distributed nature of the system, sessions from a particular MME are offloaded on different session managers. Therefore, all session managers are notified when a session is offloaded. Also, the functionality is designed to not allow all session managers to message the DEMUX manager.

- In DDN Delay feature, DDN delay timer support is at Session Management Function.
- When first data packet arrives, Sx/N4 Report message is initiated but DDN message is initiated from Session Management Function after the expiry of Delay timer.
- DDN Delay feature is a peer level feature and so, it is applied for all the session on that peer from where the DDN Delay value is received.
- In case a previous delay value was received from a peer and it is absent in the current message, the delay value will be considered as 0.

Session Recovery and ICSR is supported for DDNs.

5G SMF Calls

Downlink Data Notification - 5G UE

When UE turns to Idle state in the 5G call mode, SMF sends Sx_Modify_Request with FAR Apply Action set to BUFFER value. For every QFI (default/dedicated), the FAR Apply Action is set to BUFFER value.

Once the Downlink packet is received from the server, UPF sends the Sx_Report_Request with Downlink Data Notification with Rule Id (PDR Id)/QFI as per packet rulematch. UPF continues buffering packets until the packet limit reaches 5 for every FAR.

When UE is active, Sx_Modify_Request reaches UPF with FAR Apply Action set to FWD (Forward) and TEID (Tunnel Endpoint Identifier). UPF debuffers the packets and sends them to UE as FIFO. For each packet, rule match will take place after the debuffering process.

DDN Throttling Support

Too many DDN requests toward MME from SGW-C could lead to processing overload at MME. To reduce this load, MME dynamically requests SGW-C to reduce a certain percentage of DDN messages sent toward it for a given period time.

For DDN throttling, S-GW is required to drop a given percentage of DDNs over a given period of time. S-GW implements this functionality using a probabilistic algorithm at each session manager.

Whereas the conventional implementation of DDN throttling requires each session manager to share its list of pending DDNs for low-priority bearers with a central entity that would then calculate the net load of pending DDNs and then decide how many DDNs each session manager would have to drop. This implementation would require buffering of DDN messages at session manager. Also, due to distributed processing nature of

software subsystem in chassis, it would require considerable amount of messaging between the session managers and the central entity (demuxmgr in case of Boxer) at regular intervals.

Implementing a probabilistic algorithm removes the need for buffering at session manager and also messaging with demuxmgr. Accuracy of probabilistic algorithm increase with increasing low ARP priority paging load at session manager. Even with lower paging load, accuracy would be fairly close to the throttling factor provided.

For non-release 10 compliant MME, SGW_C provides option to enable throttling through the CLI.

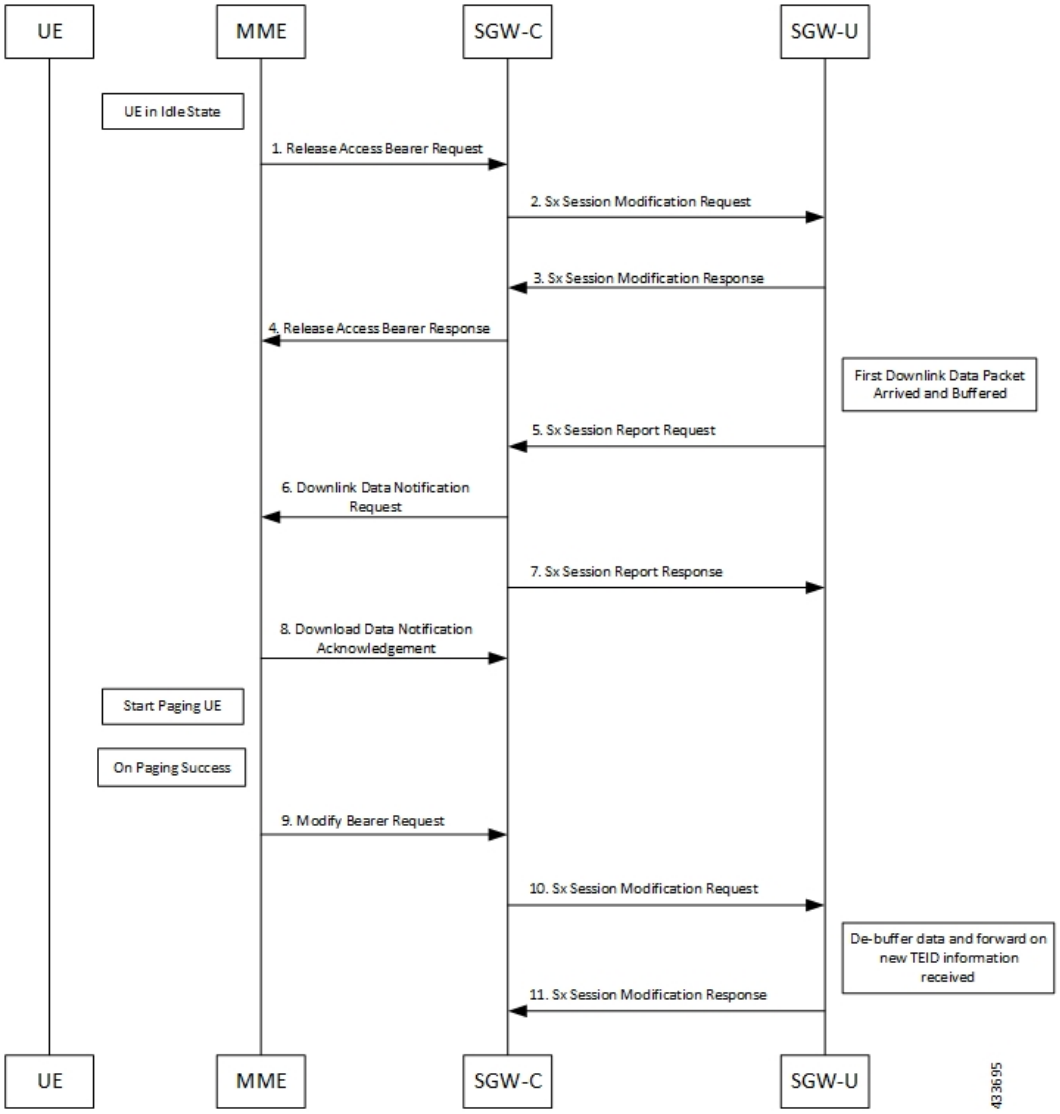
Threshold ARP values for low-priority bearer must be configured through S-GW Service Configuration. For example, if configured ARP value is 9, any bearer with ARP > 9 is considered low-priority bearer. DDN throttling is enabled through this configuration. If DDN throttling is enabled through SGW service configuration, each DDN message toward MME would contain the ARP IE.

No User Connect Timer Support

- Timer is introduced when a Modify Bearer Request is not received after positive Downlink Data Notification acknowledgment.
- It is initiated at SGW-C when DDN acknowledgment is received.
- On arrival of Modify Bearer Request, SGW-C stops this timer.
- On timer expiry SGW-C informs SGW-U to drop buffered packets.

DDN Call Flows

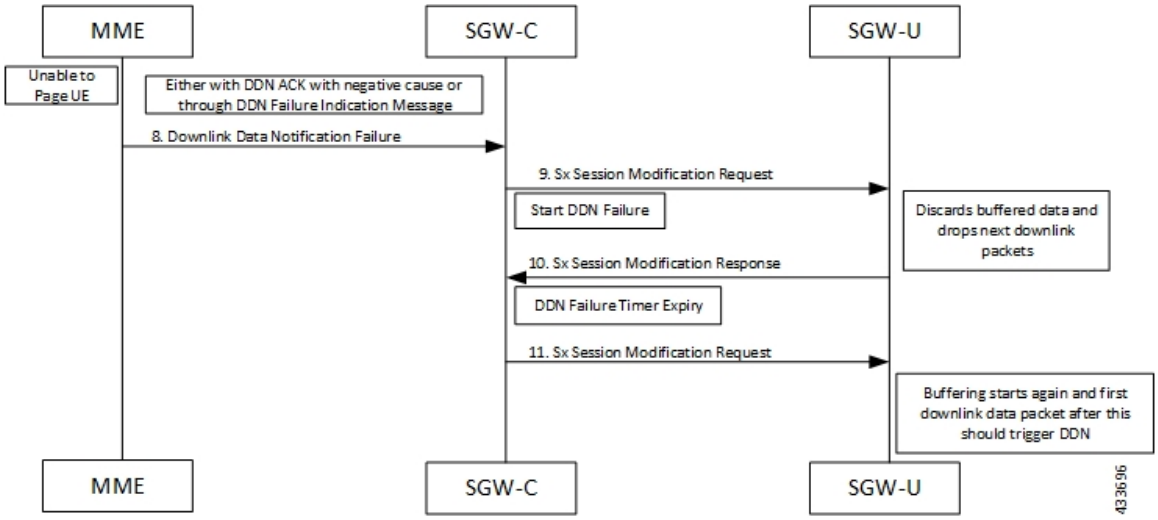
DDN Success Scenario



1. MME sends Release Access Bearer request to SGW-C to release downlink remote TEIDs of all the bearers for that UE.
2. On arrival of Release Access Bearer request, SGW-C informs the same to SGW-U by updating FAR with Apply Action as BUFFER in Sx Modification Request for all the PDNs.
3. SGW-U send Sx Modification response after applying Buffering in SGW-U for corresponding PDN.
4. SGW-C sends Release Access Bearer response to MME.
5. First Downlink data arriving in SGW-U triggers Sx Report Request (with Report Type as Downlink Data Report) towards SGW-C.

6. On arrival of Sx Report Request message, the SGW-C initiates Downlink Data Notification request message towards MME.
7. SGW-C sends Sx Report Response message towards SGW-U.
8. If MME is able to send a paging request towards UE, it sets the cause as “Request Accepted” in Downlink Data Notification Acknowledgment Message and sends it to SGW-C.
9. On successful paging, MME sends a Modify Bearer request to the S-GW with eNodeB TEIDs that sets up the S1-U connection at the SGW.
10. SGW-C sends Sx Modification request with updated FAR for new TEID information to SGW-U. SGW-U can now forward all the buffered data to UE through eNodeB.
11. SGW-U sends Sx Modification response to SGW-C.

DDN Failure Scenario



1. MME sends Release Access Bearer request to SGW-C to release downlink remote TEIDs of all the bearers for that UE.
2. On arrival of Release Access Bearer request, SGW-C informs the same to SGW-U by updating FAR with Apply Action as BUFFER in Sx Modification Request for all the PDNs.
3. SGW-U send Sx Modification response after applying Buffering in SGW-U for corresponding PDN.
4. SGW-C sends Release Access Bearer response to MME.
5. First Downlink data arriving in SGW-U triggers Sx Report Request (with Report Type as Downlink Data Report) toward SGW-C.
6. On arrival of Sx Report Request message, the SGW-C initiates Downlink Data Notification request message toward MME.
7. SGW-C sends Sx Report Response message toward SGW-U.
8. If MME is not able to page UE then it can reject Downlink Data Notification Request with relevant cause.

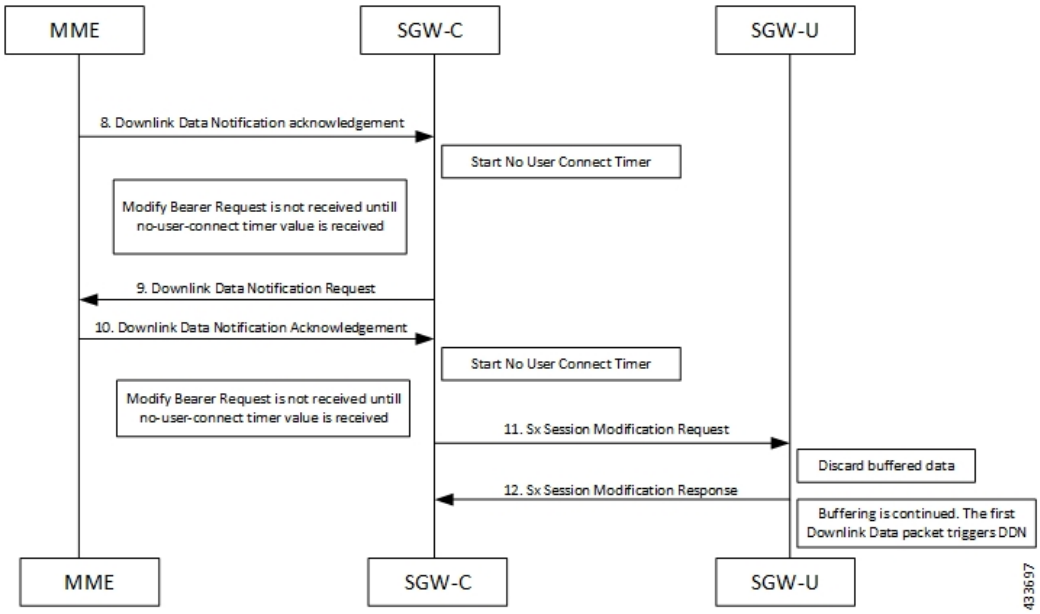
OR

If MME accepts Downlink Data Notification Request. But later sends Downlink Data Notification Failure indication in order to indicate SGW-C that the UE did not respond to paging.

9. SGW-C received DDN failure and hence to stop sending next DDN immediately , SGW-C starts DDN Failure Timer. SGW-C sends Sx Modification Request with DROBU flag to discard buffered packets and Apply Action as DROP to drop subsequent packets.
10. SGW-U sends Sx Modification Response to SGW-C.
11. On DDN Failure Timer Expiry SGW-C initiates Sx Modification with Apply Action as BUFFER in order to start buffering again.

Further steps are continued from Step 3 in the DDN Success Scenario, on page 203 call flow.

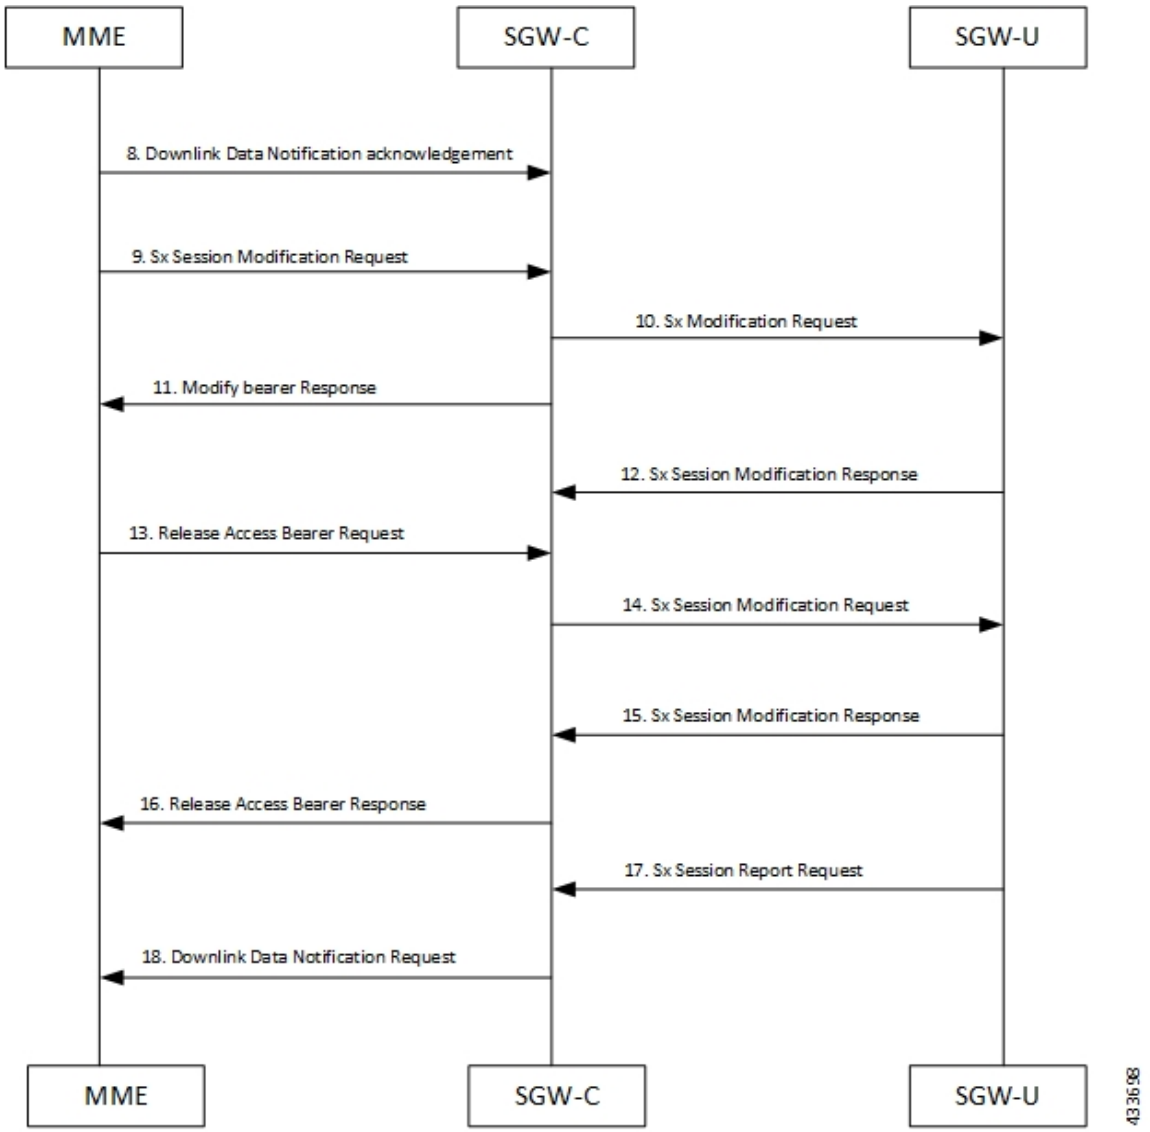
No User Connect Timer Support



1. MME sends Release Access Bearer request to SGW-C to release downlink remote TEIDs of all the bearers for that UE.
2. On arrival of Release Access Bearer request, SGW-C informs the same to SGW-U by updating FAR with Apply Action as BUFFER in Sx Modification Request for all the PDNs.
3. SGW-U send Sx Modification response after applying Buffering in SGW-U for corresponding PDN.
4. SGW-C sends Release Access Bearer response to MME.
5. First Downlink data arriving in SGW-U triggers Sx Report Request (with Report Type as Downlink Data Report) towards SGW-C.
6. On arrival of Sx Report Request message, the SGW-C initiates Downlink Data Notification request message towards MME.
7. SGW-C sends Sx Report Response message towards SGW-U.

8. Downlink Data Notification Acknowledgment is received from MME.SGW-C starts no-user-connect timer.
9. If the Modify Bearer request with eNodeB TEID information is not received and no-user-connect timer expires, SGW-C sends Downlink Data Notification again.
10. Downlink Data Notification Acknowledgment is received from MME. SGW-C initiates the no-user-connect timer again.
11. SGW-C initiates Sx Session Modification request towards SGW-U with DROBU flag set in the message. On receiving this flag SGW-U drops the buffered data. New data will be buffered, and the subsequent first packet initiates a Sx Report message for initiating Downlink Data Notification message.
12. SGW-U sends Sx Modification Response.

DDN Delay Timer



1. MME sends Release Access Bearer request to SGW-C to release downlink remote TEIDs of all the bearers for that UE.
2. On arrival of Release Access Bearer request, SGW-C informs the same to SGW-U by updating FAR with Apply Action as BUFFER in Sx Modification Request for all the PDNs.
3. SGW-U send Sx Modification response after applying Buffering in SGW-U for corresponding PDN.
4. SGW-C sends Release Access Bearer response to MME.
5. First Downlink data arriving in SGW-U triggers Sx Report Request (with Report Type as Downlink Data Report) towards SGW-C.
6. On arrival of Sx Report Request message, the SGW-C initiates Downlink Data Notification request message towards MME.
7. SGW-C sends Sx Report Response message towards SGW-U.
8. Downlink Data Notification Acknowledgment is received from MME with DDN Delay Timer value. This timer value will be saved for this peer , and now onwards every Downlink Data notification that we initiate should be after this delay for that peer.
9. On success paging, MME sends a Modify bearer request to the SGW with eNodeB TEIDs that sets up the S1-U connection at the SGW.
10. SGW-C sends Sx Modification Request with updated FAR for new TEID information to SGW-U. SGW-U can now forward all the buffered data to UE via eNodeB.
11. SGW-C sends Modify Bearer Response to MME.
12. SGW-U sends Sx Modification Response to SGW-C.
13. MME sends Release Access Bearer Request to SGW-C to release downlink remote TEIDs of all the bearers for that UE.
14. On arrival of Release Access Bearer Request, SGW-C inform the same to SGW-U via updating FAR with Apply Action as BUFFER in Sx Modification Request for all the PDNs.
15. SGW-U send Sx Modification Response after applying Buffering in SGW-U for corresponding PDN.
16. SGW-C sends Release Access Bearer Response to MME.
17. First Downlink data arriving in SGW-U triggers Sx Report Request (with Report Type as Downlink Data Report) towards SGW-C.
18. On arrival of Sx Report Request message, SGW-C starts DDN Delay Timer. On DDN Delay timer expiry SGW-C Initiates Downlink Data Notification message towards MME.

Sx Interface

Sx Session Level Reporting Procedure

Detection of first Downlink Data for Idle-Mode UE (by SAEGW-U):

When SAEGW-U receives the downlink packet but no S1-bearer for transmission and the buffering is performed by SAEGW-U, it reports the detection of first downlink data to SAEGW-C, for the purpose of paging the UE.

PFCP Session Report Request

The PFCP Session Report Request is sent over the Sxab interface by the User Plane function to report information related to a PFCP session to the Control Plane function.

Information elements	P	Condition / Comment	Appl.				IE Type
			Sxa	Sxb	Sxc	N4	
Report Type	M	This IE shall indicate the type of the report.	X	X	X	X	Report Type
Downlink Data Report	C	This IE shall be present if the Report Type indicates a Downlink Data Report.	X	-	-	X	Downlink Data Report

Downlink Data Report IE within PFCP Session Report Request

The Downlink Data Report grouped IE is encoded as shown in the following table.

Octet 1 and 2		Downlink Data Report IE Type = 83 (decimal)					
Octets 3 and 4		Length = n					
Information elements	P	Condition / Comment	Appl.				IE Type
			Sxa	Sxb	Sxc	N4	

PDR ID	M	This IE shall identify the PDR for which downlink data packets have been received at the UP function. More than one IE with this type may be included to represent multiple PDRs having received downlink data packets.	X	-	-	X	PDR ID

Notification to User Plane Function for DDN Failure

The Control Plane function notifies User Plane function for any failure so that buffered packets can be dropped and DDN related flags can be reset through DROBU flag in PFCP Sx Modification message.

PFCPSMReq-Flags	C	DROBU (Drop Buffered Packets): The CP function shall set this flag if the UP function is requested to drop the packets currently buffered for this PFCP session (see NOTE 1).

Limitations

Following are the known limitations of this feature:

- SAEGW Buffering is done for five data packets per PDN session.
- DDN profile configuration is not supported.
- Support for buffered data (data packet stream) that get deleted due to Flow Idle Timeout or other cases, is not present.

DDN Throttling for non-Release 10 Compliant MME

Use the following configuration to configure DDN throttling for a non-release 10 MME:

```
configure
   context context_name
      sgw-service service_name
         ddn throttle arp-watermark arp_value [ rate-limit limit time-factor
seconds throttle-factor percent increment-factor percent [ poll-interval seconds
 ] throttle-time-sec seconds [ throttle-time-min minutes ] [
throttle-time-hour hour ] stab-time-sec seconds [ stab-time-min minutes ] [
stab-time-hour hour ]
         no ddn throttle
         end
```

NOTES:

- **rate-limit**: DDN permitted per second.
- **time-factor**: Time period in seconds over which SGW makes throttling decision (valid range 1-300 seconds).
- **arp-value**: Valid ARP value between 1 and 15. All the packets which have ARP greater than the configured values are throttled as per the throttling factor.
- **throttling-factor**: Percentage of DDN to be dropped upon detecting DDN surge (valid range between 1-100).
- **throttling-time-sec**: Time period in seconds over which DDN is throttled at SGW (valid range between 0-59 seconds).
- **throttling-time-min**: Time period in minutes over which DDN is throttled at SGW (valid range between 0-59 minutes).
- **throttling-time-hour**: Time period in hours over which DDN is throttled at SGW (valid range between 0-310 hours).
- **increment-factor**: Percentage value by which throttling factor is incremented dynamically, if existing throttling factor is insufficient to curb the DDN surge.
- **poll-interval**: Time in seconds (optional argument, default value = 1 second, poll interval < time-factor)
- **stab-time-sec/min/hours**: Stabilization time factor, time period over which if DDN rate returns to normal, then throttling need not be applied over the entire throttling time period.

DDN throttling for non-Release-10 compliant MME makes use of existing Release-10 throttling implementation at SGW. By providing a configuration mechanism for SGW service, operator can still apply ddn throttling without needing any feedback from MME. Some salient points of this feature are described below:

1. The CLI configuration is applied per MME/S4-SGSN. Throttling parameters are tracked independently per MME/S4-SGSN.
2. On configuring this feature through CLI, demuxmgr polls each sessmgr for number of DDNs sent. By default, polling is done every second. This time interval can be changed by configuring the poll-interval time. Greater the poll interval time, lesser the number of internal messages within the chassis. However, it would take longer to detect a DDN surge.
3. By configuring time-factor, operator can specify the time interval for S-GW to apply throttling, if needed. It allows for some surge of DDNs if the net DDN rate is within specified limit over time-factor time interval. For example, time-factor= 10 seconds, ddn rate = 1000, poll interval = 2 seconds. Demux would poll each sessmgr every 2 seconds. Acceptable DDN rate limit is 1000*10 = 10000 DDNs every 10

seconds. Say after 2 seconds, 4000 DDNs were sent, in that case S-GW wouldn't apply throttling until rate limit of 10000 DDNs is crossed within time period of 10 seconds. This allows for intermittent bursts of DDNs.

4. DDN rate limit is configured through CLI. For example, if DDN rate limit is 1000 and poll interval = 1 second, time-factor = 5 seconds, then acceptable rate limit is 5000 DDNs over 5 seconds. If the number of DDNs sent by S-GW is greater than 5000 after 5 seconds, demuxmgr would ask all sessmgrs to initiate throttling.
5. Percentage of DDNs to be throttled is configured through throttling-factor.
6. Operator can specify increment-factor to increment throttling factor if the existing throttling factor is insufficient to curb the DDN surge. For example, if throttling-factor = 10%, ddn-rate = 1000, increment-factor=10%. Once throttling is applied, S-GW drops ~10% DDNs. However, if DDN rate is still greater than 1000, S-GW would increase throttling-factor to 20%. If this is still not sufficient, it would be incremented to 30%. After incrementing throttling factor, if number of DDNs dropped are greater than expected, throttling-factor would then be decrement by increment-factor. For example, in this case, after increasing throttling factor to 30%, if DDNs sent is less than 1000 per second (taking time-factor and poll-interval into consideration), throttling factor would be decremented to 20. The cap for decrementing throttling-factor would be the configured value (10% in this case).
7. Operator can configure the time duration for which throttling is applicable at S-GW. This could be a large value in order of days (for example: 10 days or 240 hours). The operator has an option to stop throttling if DDN rate is well under control by configuring stabilization time factor. In such a case, DDNs won't be needlessly dropped. For example, throttling-time =10 days, stab-time = 8 hours. After S-GW starts DDN throttling, in a time span of 8 hours, DDNs sent + DDNs dropped < ddn-rate * 8 hours, throttling would be stopped.

DDN Throttling for Release 10 Compliant MME

DDN throttling is enabled through Call Control Profile by providing the ARP value. For example, if the ARP value provided is 10, then all bearers with ARP value between 10-15 are treated as low priority bearers and are given throttling treatment. Throttling would not be enabled if ARP value is not provided through S-GW service configuration. Also, ARP IE in DDN message towards MME would not be included unless DDN throttling is configured using S-GW service. If MME is Release 10 compliant, the user need not configure the duration value as the DDN Acknowledgment would have the throttling IE. Otherwise, throttling can be enabled at S-GW by setting the duration value. If it's set to 0, S-GW would apply throttling recurringly. To enable throttling only for a given duration of time (in non Rel-10 compliant MME), user needs to set the value in hours and minutes. From the time of configuration, throttling would be applied at S-GW until the timer duration expires. For example, if user sets hours = 10, minutes = 30, S-GW would apply throttling for next 10 hours 30 minutes.

On re-configuration, all the parameters will be set with new values, but they will be applicable only from the next recalibration except from polling time and time factor.

Use the following configuration to configure DDN throttling for release 10 MME:

```
configure
  context context_name
    sgw-service service_name
      [ no ] ddn throttle arp-watermark arp_value
      end
```

NOTES:

- **arp-value**: Valid ARP value between 1 and 15. All the packets which have ARP greater than the configured values will be throttled as per the throttling factor.

Idle Timer for SAE-GW Sessions

An Idle Timer is supported to identify and remove idle sessions that occur in the SAE-GW.

A session becomes idle in some cases where the session is removed from other network nodes, but due to a technical mishap the session could still remain on the SAE-GW leading to resources being held by these idle sessions.

The Idle Timer, once configured, removes those sessions that remain idle for longer than the configured time limit effectively utilizing the system capacity.



Important This feature is currently restricted to Pure-P and Collapsed Call.

Limitations

The Idle Timer feature does not support recovery of Idle Timer in case of redundancy events.

Configuring Idle Timer for SAE-GW Sessions

The Idle Timer is configurable at APN level.

Use the following commands to configure the idle timer for SAE-GW sessions:

```
configure
  context context_name
    apn apn_name
      timeout idle timeout_value
      no timeout idle
      default timeout idle
      end
```

- **no**: Disables the idle timer configuration.
- **default**: Configures the default value for subscriber's time out settings. The default idle timeout value is 0.
- **idle** *timeout_value*: Designates the maximum duration a session can remain idle, in seconds, before system automatically terminates the session. Must be followed by number of seconds between 0 and 4294967295. Zero indicates function is disabled.

S-GW Session Idle Timeout

This chapter describes the Idle Timeout Handling feature for S-GW sessions. On the ASR5500 platform, subscriber session is represented by call-line. The S-GW product call-line interfaces to its peers through MME/S4-SGSN on S11/S4 and P-GW on S5/S8. In some scenarios, peer sessions are deleted by respective peers, S-GW does not receive or miss deletion messages, and as a result S-GW session remains idle. Such idle or stale sessions are counted toward valid call-lines in system for effectively consuming resources and causing capacity reduction. In such cases, S-GW triggers to get the new subscriber session, which results in the removal of old session for same subscriber. The Idle Timeout Handling support enables the identification of such sessions and initiates deletion to release the resources.

The following points describe the idle timeout handling for S-GW sessions:

- The subscriber session is idle when there is no data traffic activity for the subscriber. The session manager keeps track of the call-line state, when no data traffic is recorded for call-line, such sessions are moved to idle state.
- Session which is idle for defined timeframe referred as idle timeout is considered for idle timeout handling. In idle timeout session, S-GW initiates the deletion of session toward its peers.
- Idle timeout is configured in seconds depending on the network requirements. The timeout range is 1-4294967295 seconds.
- The idle timeout configuration is applicable on S-GW service level for enabling the idle timeout handling for set of subscribers handled by that service.

Configuring Session Idle Timeout

The session idle timer for S-GW sessions is configurable from S-GW service.

To configure Session Idle Timeout for S-GW, use the following configuration:

```
configure
   context context_name
      sgw-service service_name
         [ no | default] timeout idle timeout_duration
         end
```

NOTES:

- **timeout idle** *timeout_duration*: Specifies the maximum duration a session can remain idle for, in seconds, before the system automatically terminates the session. *timeout_duration* must be an integer in the range of 1-4294967295. 0 disables the feature. By default, it is disabled for the S-GW service.

Show Commands Input and/or Outputs

This section provides information regarding show commands and their outputs in support of the feature.

show subscribers user-plane-only full all

The output of this command displays the following fields in support of this feature:

- buffered pkts
- buffered bytes
- buffer overflow drop pkts
- buffer overflow drop bytes

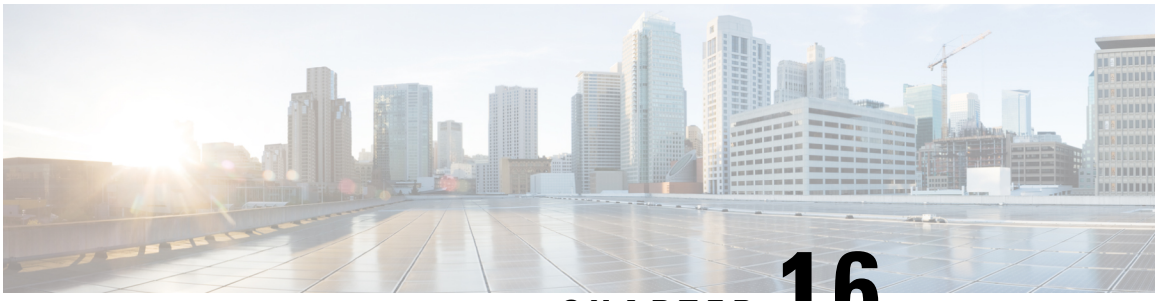
show subscribers user-plane-only full callid <call_id>

Use the following configuration to check the buffering per subscriber:

- DDN buffered packets
- DDN buffered bytes
- DDN buffer overflow drop packets
- DDN buffer overflow drop bytes

When the buffered packets are debuffered the state changes back to zero.

Each time the packet limit reaches 5, the additional packets get dropped as overflow drops.



CHAPTER 16

DSCP Markings For Collapse Calls

- Feature Summary and Revision History, on page 215
- Feature Description, on page 216
- How It Works, on page 216
- Configuration, on page 219
- Monitoring and Troubleshooting, on page 220

Feature Summary and Revision History

Table 44: Summary Data

Applicable Product (s) or Functional Area	5G-UPF
Applicable Platforms	VPC-SI SMI
Feature Default Setting	Disabled – Configuration Required
Related Changes in this Release	Not Applicable
Related Documentation	Not Applicable

Table 45: Revision History

Revision Details	Release
Provided updated output for **show subscribers user-plane-only full all** and **show user-plane-service statistics qos-group sessmgr all** CLI commands.	2022.04.0
First introduced	2021.02.0

Feature Description

Currently, QCI-based DSCP markings are applicable for Pure-S and Pure-P calls. The DSCP markings are based on QCI-QOS-Mapping associated with respective S-GW service or P-GW service. For collapse calls QCI-QOS-Mapping associated with PGW-service is applicable. This feature helps to apply the DSCP markings for collapse calls based on associated S-GW and P-GW services for uplink and downlink traffic. For uplink traffic, DSCP markings associated with logical P-GW service are applicable. For downlink traffic, DSCP markings associated with logical S-GW service are applicable. The DSCP markings are present in IP header of data traffic as a part of GTP-U header and Inner IP. There's option to enabled or disable this functionality by CLI configuration. When you enable the feature, then only the new functionality is applicable otherwise existing functionality also works. By default, this feature is disabled, so that there's no impact on customers who upgrades to this feature.

DSCP Markings for 5G Calls

The QCI/QOS mapping table in SMF drives the DSCP values being put on the packets. It's similar configuration as in PGW-CP. The DSCP values are sent as part of FAR during call establishment (Sx establishment request/Sx modification request). The DSCP values are being applied to inner packet and outer GTPU packet.

DSCP Markings for 4G Collapsed Datapath Calls

For collapsed Datapath calls:

- On Sxa leg, DSCP marking is sent in FAR from SGWC/Cn-SGWC.
- On N4 leg, DSCP marking is sent in FAR from SMF.
- For uplink packets, UPF applies the DSCP marking as per SMF configuration.
- For downlink packets, UPF applies the DSCP marking as per SGWC/Cn-SGWC configuration in GTPU header.
- For inner packet, DSCP marking is as per SMF configuration.

The SMF configuration/logic of sending the DSCP marking is similar as PGW-C. In case of the DSCP marking for the ECS charging action, priority is given to the charging action configuration for inner packet DSCP marking.

How It Works

Following are the steps that describe the DSCP markings for the collapse calls.

- In case of Collapse call:
  - For ACCESS side, QCI-QOS mapping table associated with SGW-service is used.
  - For CORE side, QCI-QOS mapping table associated with PGW-service is used.
- The preceding conditions apply once you enable the feature, otherwise QCI-QOS mapping table associated with PGW-service is applicable for both sides.

- APN associated QCI-QOS mapping table is preferred over the P-GW service QCI-QOS mapping table.
- APN-Profile associated QCI-QOS mapping table is preferred over SGW-Service QCI-QOS mapping table for ACCESS side DSCP markings.
- In case, only P-GW service has QCI-QOS mapping table configuration, then these DSCP markings is applicable on both ACCESS & CORE side for collapse call.
- In case only S-GW service has QCI-QOS mapping table configuration then these DSCP markings is applicable on ACCESS side for collapse call.
- There's a new configurable parameter inside the SAEGW service which indicates whether the feature is enabled or disable.
- For Pure-P to Collapse HO and conversely, transport layer markings are updated in FAR as a part of Sx Modify request.
- Layer2 markings are also modified based on QCI-QOS mapping table picked for ACCESS and CORE side.
- DSCP markings continue to apply on existing bearers post session recovery.
- DSCP markings continue for the bearers on standby chassis once it switches to active mode.

SessMgr SMF Changes

DSCP markings for Uplink/CORE and Downlink/ACCESS are present at bearer level inside sessmgr_sub_session_t → sessmgr_qci_tab_t.

User datagram DSCP markings are updated in IP header of inner packet. That is, packet sent from UE to Internet and the opposite way.

Encaps header DSCP markings are updated in IP header of outer IP layer having GTP-U header (Outer header).

DSCP markings are sent from SMF to UPF inside FAR IE as follows:

- Transport Level Marking - The DSCP markings is configured in encaps header for ACCESS side and User-datagram on CORE side for collapse call.
- Transport Level Marking Options—Includes two options and are applicable only for outer header:
  - Copy-inner: Copy the inner packets markings to outer header
  - Copy-outer: Relay the DSCP markings for outer header

Inner Packet Marking—DSCP markings is configured in user datagram for ACCESS side. For CORE side, it's N/A for collapse call.

Logic to fetch the DSCP marking has changed for collapse call:

- Fetch the DSCP markings based on QCI and "qrp_pl" for session from the associated SGW Service for ACCESS/downlink side.
- Fetch the DSCP markings based on QCI and "qrp_pl" for session from the associated PGW Service for CORE/uplink side.
- For ACCESS/downlink side, QCI-QOS-mapping table associated with APN-profile takes preference over SGW Service QCI-QOS-mapping table.

- For CORE/uplink side QCI-QOS-mapping table associated with APN config takes preference over PGW Service QCI-QOS-mapping table.
- In case SGW Service QCI-QOS-mapping table isn't configured, then PGW Service QCI-QOS-mapping table is applicable on both ACCESS/CORE side.
- In case PGW Service QCI-QOS-mapping table isn't configured, then SGW Service QCI-QOS-mapping table is applicable on ACCESS side and no DSCP markings are applicable on CORE side.
- DSCP markings are updated on UPF in create/update FAR sent as a part of Sx/N4 Establishment/Modification request from SMF to UPF.
- Update the TLM, IPM, and TLMO in case of HO from Pure-P to Collapse and vice versa in Sx/N4 Modification request as a part of Update FAR IE.
- Update the Layer2 markings in case of HO from Pure-P to Collapse and vice versa in Sx/N4 Modification request as a part of Update FAR IE.

Following table depicts the various possible config combinations and outcome for DSCP markings to be applied on ACCESS and CORE side for COLLAPSE call:

S. No.	Feature Enable / Disable	PGW Service QOS-QCI Table Configured (Q1)	SGW Service QOS-QCI Table Configured (Q2)	APN QOS-QCI Table Configured (Q3)	APN-Profile QOS-QCI Table Configured (Q4)	ACCESS/Downlink DSCP Markings for Collapse Call	CORE/Uplink DSCP Markings for Collapse Call
1	ENABLE	YES	YES	YES	YES	Q4 (APN-Profile)	Q3(APN)
2	ENABLE	YES	YES	YES	NO	Q2 (SGW-Service)	Q3(APN)
3	ENABLE	YES	YES	NO	YES	Q4 (APN-Profile)	Q1 (PGW-service)
4	ENABLE	YES	YES	NO	NO	Q2 (SGW-Service)	Q1 (PGW-service)
5	ENABLE	YES	NO	YES	YES	Q4 (APN-Profile)	Q3(APN)
6	ENABLE	YES	NO	YES	NO	Q3(APN)	Q3(APN)
7	ENABLE	YES	NO	NO	YES	Q4 (APN-Profile)	Q1 (PGW-service)
8	ENABLE	YES	NO	NO	NO	Q1 (PGW-service)	Q1 (PGW-service)
9	ENABLE	NO	YES	YES	YES	Q4 (APN-Profile)	Q3(APN)
10	ENABLE	NO	YES	YES	NO	Q2 (SGW-Service)	Q3(APN)
11	ENABLE	NO	YES	NO	YES	Q4 (APN-Profile)	N/A (NO DSCP)
12	ENABLE	NO	YES	NO	NO	Q2 (SGW-Service)	N/A (NO DSCP)
13	ENABLE	NO	NO	YES	YES	Q4 (APN-Profile)	Q3(APN)
14	ENABLE	NO	NO	YES	NO	Q3(APN)	Q3(APN)
15	ENABLE	NO	NO	NO	YES	Q4 (APN-Profile)	N/A (NO DSCP)
16	ENABLE	NO	NO	NO	NO	N/A (NO DSCP)	N/A (NO DSCP)
17	DISABLE	YES	YES	YES	YES	Q3(APN)	Q3(APN)

S. No.	Feature Enable / Disable	PGW Service QOS-QCI Table Configured (Q1)	SGW Service QOS-QCI Table Configured (Q2)	APN QOS-QCI Table Configured (Q3)	APN-Profile QOS-QCI Table Configured (Q4)	ACCESS/Downlink DSCP Markings for Collapse Call	CORE/Uplink DSCP Markings for Collapse Call
18	DISABLE	YES	YES	YES	NO	Q3(APN)	Q3(APN)
19	DISABLE	YES	YES	NO	YES	Q1 (PGW-service)	Q1 (PGW-service)
20	DISABLE	YES	YES	NO	NO	Q1 (PGW-service)	Q1 (PGW-service)
21	DISABLE	YES	NO	YES	YES	Q3(APN)	Q3(APN)
22	DISABLE	YES	NO	YES	NO	Q3(APN)	Q3(APN)
23	DISABLE	YES	NO	NO	YES	Q1 (PGW-service)	Q1(PGW-service)
24	DISABLE	YES	NO	NO	NO	Q1 (PGW-service)	Q1 (PGW-service)
25	DISABLE	NO	YES	YES	YES	Q3(APN)	Q3(APN)
26	DISABLE	NO	YES	YES	NO	Q3(APN)	Q3(APN)
27	DISABLE	NO	YES	NO	YES	N/A (NO DSCP)	N/A (NO DSCP)
28	DISABLE	NO	YES	NO	NO	N/A (NO DSCP)	N/A (NO DSCP)
29	DISABLE	NO	NO	YES	YES	Q3(APN)	Q3(APN)
30	DISABLE	NO	NO	YES	NO	Q3(APN)	Q3(APN)
31	DISABLE	NO	NO	NO	YES	N/A (NO DSCP)	N/A (NO DSCP)
32	DISABLE	NO	NO	NO	NO	N/A (NO DSCP)	N/A (NO DSCP)

Configuration

Configure the following command inside the SAEGW service to enable/disable this feature.

```
configure
  context egress
    saegw-service saegw_service_name
      downlink-dscp-per-call-type [ enabled | disabled ]
      end
```

Note For collapsed call, when you enable the feature, use the S-GW service QCI-QOS mapping DSCP markings for downlink. By default, the downlink-DSCP-per-call-type is Disabled.

Monitoring and Troubleshooting

This section provides information on CLI commands that are available for monitoring and troubleshooting for DSCP markings for collapse calls.

Show Commands Outputs

This section provides information about show CLI commands that are available in support of DSCP markings for collapse calls.

show saegw-service all

This show command is to check if the feature is enabled or Disabled.

```
Service name : SAEGW11
Service-Id : 47
Context : EPC1
Status : STARTED
sgw-service : SGW11
pgw-service : PGW11
sx-service : SX11C
User Plane Tunnel GTPU Service : SAEGW11SXU
Newcall policy : n/a
downlink-dscp-per-call-type : enabled
CUPS Enabled : Yes
Service name : SAEGW21
Service-Id : 25
Context : EPC2
Status : STARTED
sgw-service : SGW21
pgw-service : PGW21
sx-service : SX21C
User Plane Tunnel GTPU Service : SAEGW21SXU
Newcall policy : n/a
downlink-dscp-per-call-type : disabled
CUPS Enabled : Yes
```

show subscribers user-plane-only callid <call_id> far full all

Use this User Plane CLIs to validate the Transport level marking options and inner packet markings for UPLINK/DOWNLINK FAR.

show subscribers user-plane-only full all

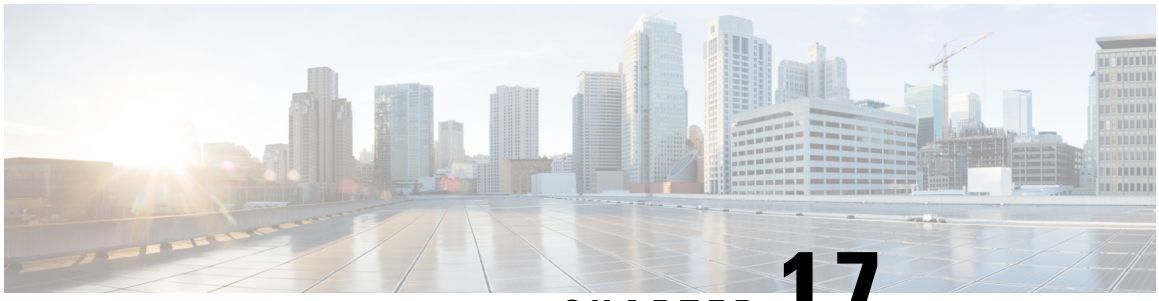
Use this User Plane CLI to see the number of TOS marked packets for U/L and D/L.

```
ToS marked Uplink Pkts: 0
ToS marked Downlink Pkts: 0
```

show user-plane-service statistics qos-group sessmgr all

Use this User Plane CLI to see the statistics for U/L and D/L packets of QoS group Sessmgr instance.

```
Uplink Packets Marked: 0
Downlink Packets Marked: 0
```



CHAPTER 17

Dynamic and Static PCC Rules

This chapter covers the following topics:

- Feature Summary and Revision History, on page 221
- Feature Description, on page 222
- Provisioning of Predefined PCC Rules, on page 222
- Dynamic PCC Rules Support, on page 223
- Policing, on page 224
- Bandwidth Policy Configuration Limits, on page 226
- Rate Limiting for Static and Predefined Rules, on page 226
- Rate Limiting for Dynamic Rules, on page 227
- Standards Compliance, on page 228
- Configuring the URR IDs, on page 228
- Threshold Configuration, on page 229

Feature Summary and Revision History

Summary Data

Table 46: Summary Data

Applicable Product(s) or Functional Area	5G-UPF
Applicable Platform(s)	VPC-SI SMI
Feature Default Setting	Enabled – Always-on
Related Changes in this Release	Not Applicable
Related Documentation	*UCC 5G UPF Configuration and Administration Guide*

Revision History

Table 47: Revision History

Revision Details	Release
Added support for: • Flow-level policing • Increase in the the maximum number of groups per bandwidth policy	2021.01.0
First introduced.	2020.02.0

Feature Description

Dynamic PCC rules are provisioned by the PCF to the PCEF through the HTTP interface and may be either predefined/static or dynamically generated in the PCF. Dynamic PCC rules can be installed, modified and removed at any time.

Predefined PCC rules are configured in the PCEF and can be activated or deactivated by the PCF or by the PCEF at any time. Static PCC rules within the PCEF may be grouped allowing the PCF to dynamically activate a set of static PCC rules over the HTTP reference point. Those static PCC rules to be locally activated by the PCEF are not explicitly known in the PCF, but the PCF simply knows identifiers of static PCC rules to be activated from the PCF.

How it Works

Predefined PCC Rules Support

Config URR IDs are applicable for static rules and also predefined rules. When a subscriber call comes up, it traverses the static rules in rule base. The subscriber primary URR list with bucket IDs as key updates the corresponding URR buckets for the various interfaces with the charging action configuration. For dynamic rules and predefined rules, URR ID list in PDR creates the URR buckets on the User Plane.

Following are the ecosystem changes to support Cisco SMF and UPF to work independently for Charging Action (vendor agnostic way) to work:

- Configurable "Config URR IDs" at UPF
- UPF to enable the local configuration for thresholds

Provisioning of Predefined PCC Rules

Predefined PCC rule is preconfigured in the SMF (for 5GC). Predefined PCC rules can be activated or deactivated by the PCF at any time. The Predefined PCC rules may be grouped allowing the PCF to dynamically activate a set of PCC rules. The SMF may enforce an activated predefined PCC rule by the PCF in the UPF by:

- Determining the service data filters or application IDs referred by the activated predefined PCC rule(s) and the corresponding QoS and charging control information respectively.
- Creating the necessary PDR(s) to identify the service data flow(s), application(s) that the predefined PCC or ADC rule refer to, if not already existing.
- Creating the necessary QER for the QoS enforcement at service data flow or application-level accordingly.
- Creating the necessary FAR if a new FAR must be created as result of QoS flow binding and QoS control for forwarding the detected service data flow or application traffic, or to redirect or to apply traffic steering control if included in the predefined PCC rule.
- Creating the necessary URR(s) for each monitoring key, charging key, combination of charging key and service ID, or combination of charging key, sponsor ID and Application Service Provider ID if included in the predefined PCC rule.

And, later by:

- Associating the created URR(s) to the newly created PDR(s).
- Associating the existing FAR or the new FAR to the newly created PDR(s).

Optionally, the traffic handling policies common to many PFCP sessions (that is, predefined QER(s)/FAR(s)/URR(s)) can be configured in the UPF. The SMF activates these traffic handling policies by including the Activate Predefined Rules IE within one of the following:

- The Create PDR IE in an PFCP Session Establishment Request
- The Create PDR IE in an PFCP Session Modification Request

For traffic matching PDR(s) associated with the activated predefined rules, the UPF enforces the rules. For example, the UPF generates Usage Report(s) and sends it to the SMF, for URR, and the SMF handles the usage reports.

The URR IDs used in reports triggered by a predefined rule in UPF are also preconfigured at the SMF.

Note The UPF does not support the deactivation of predefined rules.

Dynamic PCC Rules Support

For dynamic PCC rules multiple flows are supported on per Packet Forwarding Control Protocol (PFCP) session:

- The 5G QoS model allows classification and differentiation of specific services based on subscription-related and invocation-related priority mechanisms. These mechanisms provide abilities such as invoking, modifying, maintaining, and releasing QoS Flows with priority, and delivering QoS Flow packets according to the QoS characteristics under network congestion conditions.
- The 5G QoS model is based on QoS Flows. The 5G QoS model supports both QoS Flows that require guaranteed flow bit rate (GBR QoS Flows) and QoS Flows that do not require guaranteed flow bit rate (Non-GBR QoS Flows).

- The QoS Flow is the finest granularity of QoS differentiation in the PDU session. A QoS Flow ID (QFI) is used to identify a QoS Flow in the 5G System. User Plane traffic with the same QFI within a PDU session receives the same traffic forwarding treatment (Example - scheduling, admission threshold).
- Within the 5GS, a QoS Flow associated with the default QoS rule is required to be established for a PDU session and remains established throughout the lifetime of the PDU session. This QoS Flow must be a Non-GBR QoS Flow.
- A QoS flow is associated with QoS requirements as specified by QoS parameters and QoS characteristics. A QoS flow can either be "GBR" or "Non-GBR" depending on its QoS profile.
  - For each QoS Flow, the QoS profile includes the QoS parameters:
    - 5G QoS Identifier (5QI)
    - Allocation and Retention Priority (ARP)
  - For each GBR QoS flow only, the QoS profile must also include the QoS parameters:
    - Guaranteed Flow Bit Rate (GFBR) - UL and DL
    - Maximum Flow Bit Rate (MFBR) - UL and DL
  - In case of a GBR QoS Flow only, the QoS profile may also include one or more of the QoS parameters:
    - Notification control
    - Maximum Packet Loss Rate - UL and DL

During PDR creation or modification UPF receives the QER for QoS enforcement on flows.

The QoS enforcement rule correlation ID is assigned by the CP function to correlate QERs from multiple PFCP session contexts. For instance, the enforcement of APN-AMBR in the PGW-U is achieved by setting the same QoS enforcement rule correlation ID to the QERs from different PFCP sessions associated with all the PDRs corresponding to the non-GBR bearers of all the UE's PDN connections to the same APN. The QERs that are associated to the same QoS Enforcement Rule Correlation ID in multiple PFCP sessions will be provisioned with the same QER contents in each of these PFCP sessions. The QoS enforcement rule correlation ID is only used to enforce the APN-AMBR when the UE is in EPC, it may be provided by the CP function over N4 to the UP function for a PDU session may move to EPC in a later stage.

If the UPF receives QoS Enforcement Rule Correlation ID for 5G PFCP sessions, then it enforces it.

Policing

The policer configuration uses inputs from the session manager, these inputs are received either from PCF as AMBR or from flow-level QoS information. The values received from the PCF are always accepted for session-level AMBR policing. However, the flow-level policing is prioritized, if available, and AMBR policing is applied sequentially. That is to say, the policer engine applies the hierarchical policing—first the flow-level/rule bandwidth limiting and then the session-level bandwidth limiting.

Note AMBR modifications during session run-time through RAR or CCA-U is applicable.

The input values received from the session manager are pushed into a policer configuration and a policer token bucket. For each direction - uplink or downlink, a new record is created for Policer configuration and Policer token bucket.

The Policer configuration is the reference for the policer engine, and the policer token bucket is used for calculation and restoration of values.

Currently, Policing is supported for AMBR received from PCF and rule-level QoS information for dynamic rules. For static and predefined rules, bandwidth limiting is achieved by the bandwidth policy configuration. Extended bit rates configured in bandwidth-policy configuration in Active Charging Service Configuration mode on SMF is provided to the UPF by RCM, and same is applied for policing by the UPF. An example configuration of bandwidth policy, with extended bit rate, is given below:

```
configure
  active-charging service ACS
    bandwidth-policy BWP

      flow limit-for-bandwidth id 1 group-id 2

      flow limit-for-bandwidth id 2 group-id 3
      flow limit-for-bandwidth id 100 group-id 100

     group-id 2 direction uplink peak-data-rate 256000 peak-burst-size 1000 violate-action
 discard
     group-id 3 direction downlink peak-data-rate 256000 peak-burst-size 1000 violate-action
 discard
     group-id 4 direction uplink peak-data-rate 300000 peak-burst-size 1200 violate-action
 lower-ip-precedence
     group-id 5 direction downlink peak-data-rate 300000 peak-burst-size 1200 violate-action
 lower-ip-precedence committed-data-rate 256000 committed-burst-size 1000 exceed-action
lower-ip-precedence
      group-id 100 direction downlink peak-data-rate-kbps 4294967295 peak-burst-size
4294967295 violate-action discard
     group-id 100 direction uplink peak-data-rate-kbps 4294967295 peak-burst-size 4294967295
 violate-action discard
      exit
    charging-action catchall
      flow limit-for-bandwidth id 1
      exit
    rulebase cisco
      bandwidth default-policy BWP
      exit
    end
```

Limitations

In this release, Policing has the following limitations:

- Modification of **bandwidth-policy** isn't supported.
- Interaction with other features, such as token replenishment (both APN-level and ACL-level) isn't supported.
- Currently, policer-based statistics aren't supported. You can verify bandwidth limiting using network performance monitoring tools.

Bandwidth Policy Configuration Limits

The user can configure bandwidth limits in both SMF and UPF, for both downlink and uplink packets. The configuration is done in all charging actions of predefined PCC rules, even if the bandwidth limitation configuration is the same for all charging actions.

To optimize these configurations, you must define a bandwidth ID to include all bandwidth-related configurations and associate the bandwidth ID under charging actions.

If the bandwidth value changes, the new subscribers use the configured bandwidth values while the existing subscribers continue to use the old values.

The following are the bandwidth-policy configuration limits:

- Maximum number of bandwidth policies—64
- Maximum number of groups per bandwidth policy—1000
- Maximum number of bandwidth IDs per bandwidth policy—1000
- Maximum number of groups across bandwidth policies—10000
- Maximum number of bandwidth IDs across bandwidth policies—10000

Rate Limiting for Static and Predefined Rules

For static and predefined rules, bandwidth limiting is achieved by the bandwidth policy configuration. Bandwidth Policy must be configured on SMF and UPF under Active Charging Service Configuration Mode.

The following is an example configuration of bandwidth policy with extended bit rate:

```
config
   active-charging service ACS
      bandwidth-policy BWP
         flow limit-for-bandwidth id 1 group-id 2
         flow limit-for-bandwidth id 2 group-id 3
         flow limit-for-bandwidth id 100 group-id 100
        group-id 2 direction uplink peak-data-rate 256000 peak-burst-size 1000 violate-action
 discard
         group-id 3 direction downlink peak-data-rate 256000 peak-burst-size 1000
violate-action discard
        group-id 4 direction uplink peak-data-rate 300000 peak-burst-size 1200 violate-action
 lower-ip-precedence
         group-id 5 direction downlink peak-data-rate 300000 peak-burst-size 1200
violate-action
         lower-ip-precedence committed-data-rate 256000 committed-burst-size 1000
exceed-action
         lower-ip-precedence
         group-id 100 direction downlink peak-data-rate-kbps 4294967295 peak-burst-size
4294967295 violate-action discard
         group-id 100 direction uplink peak-data-rate-kbps 4294967295 peak-burst-size
4294967295 violate-action discard
         exit
      charging-action catchall
         flow limit-for-bandwidth id 1
         exit
      rulebase cisco
```

```
bandwidth default-policy BWP
exit
end
```

Note UPF does not support the modification of bandwidth-policy configuration.

Rate Limiting for Dynamic Rules

As per 3GPP TS 29.244, the following Information Element (IE) is received from SMF for QoS enforcement in Create QER or Update QER in Session Establishment or Modification Request:

- Maximum Bitrate: This IE is present if an MBR enforcement action is applied to packets matching this PDR. When present, this IE indicates the uplink and/or downlink maximum bit rate to be enforced for packets matching the PDR. For 5GC, this IE may be set to the value of:
  - Session-AMBR—for a QER that is referenced by all the PDRs of the non-GBR QoS flows of a PDU session.
  - QoS Flow MBR—for a QER that is referenced by all the PDRs of a QoS Flow.
  - SDF MBR—for a QER that is referenced by all the PDRs of an SDF.
- Guaranteed Bitrate: This IE is present if a GBR has been authorized to packets matching this PDR. When present, this IE indicates the authorized uplink and/or downlink guaranteed bit rate. This IE may be set to the value of:
  - Aggregate GBR—for a QER that is referenced by all the PDRs of a GBR bearer
  - QoS Flow GBR—for a QER that is referenced by all the PDRs of a QoS Flow
  - SDF GBR—for a QER that is referenced by all the PDRs of an SDF
- QoS flow identifier (QFI): This IE is present if the QoS flow identifier is inserted by the UPF.
- Gate Status: This IE indicates whether the packets are allowed to be forwarded (the gate is open) or it is discarded (the gate is closed) in the uplink and/or downlink directions.
- QER Correlation ID: This IE is present if the UP function is required to correlate the QERs of several PFCP sessions, for APN-AMBR enforcement of multiple UE's PDN connections to the same APN.

Note Although it is not applicable, but if UPF receives QoS Enforcement Rule Correlation ID for 5G PFCP sessions then it enforces it.

The SMF provisions QoS enforcement in UPF by creating necessary PDRs to represent SDF, QoS Flow and session and associating respective QERs as follows:

- Creating QERs for the QoS enforcement at session level, SDF level.
- Creating QERs for the QoS enforcement of the aggregate of SDFs with the same GBR QFI.

- Associating the session level QER to all the PDRs defined for the session.
- Associating the SDF or application QER to the PDRs associated to the SDF or application.
- Associating the QER of the aggregate of SDFs to the PDRs associated to SDFs or applications that share the QER.

Standards Compliance

The N4 interface between SMF and UPF is specified in 3GPP TS 23.501 and 3GPP TS 23.502.

Configuring the URR IDs

Config URR IDs are applicable for static rules and predefined rules that are bound to the default bearer. When a subscriber call comes up, static rules in the Rulebase are traversed. The corresponding URR buckets for the various interfaces configured in the Charging Action are updated in the subscriber primary URR list with bucket IDs as key. For Dynamic rules and Predefined rules that are bound to dedicated bearers, the URR ID list in PDR is used for creating URR buckets on the UP.

Currently, the Config URR IDs are generated using running counter at SessCtrl.

To achieve the configurable Config URR IDs:

- Configuration template outside of Charging Action allows URR-ID mapping with “Rating Group” and “Service-id”.
- If a separate Rating Group (RG) is configured for Gy, then that RG is applied to the Gy bucket. If no separate RG is configured for Gy, then the same “Content-id” is applicable for all interfaces.
- “Service-id” is optional for URR-ID mapping.
- URR-ID is unique. This must be ensured through **show configuration error** CLI command, a separate validation script, or checking during config time.
- For UPF, the value for URR-ID generation is taken from the configuration. The UP receives Config URR-ID from configuration rather than PFD message. The PFD mechanism is not used between SMF and UPF.
- SMF does not support the URR-ID list configuration.
- URR-ID is generated independently on UPF and SMF for Static or Predefined rules having rating group/content-id,service-id, and so on.

The following is a sample configuration for URR-ID profile template:

```
rg <x> si <y> urr-id <abc>
rg <x1> si <y1> urr-id <abc1>
rg <x> urr-id <abc2>
```

Use the following configuration to enable Config URR ID.

```
configure
   active-charging service service_name
      urr-list list_name
```

```
        rating-group group_number { service-identifier service_number | urr-id
id_range }
         end
```

NOTES:

- **urr-list** *list_name:* Configures the active charging service URR list. *list_name* must be an alphanumeric string of 1 to 63 characters.
- **rating-group** *group_number* **:** Specifies the rating ID used in prepaid charging. *group_number* must be an integer in the range of 0 to 2147483647.
- **service-identifier** *service_number* **:** Specifies the number given to the service.
- **urr-id** *id_range* **:** Specifies the URR identifier for rating/service group. *id_range* must be an integer in the range of 1-134217727.
- The values can be changed dynamically. However, it will take effect only for new sessions.

Threshold Configuration

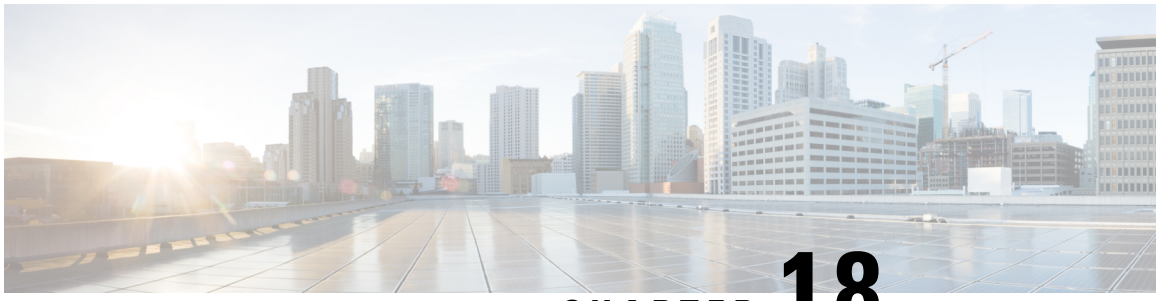
The GTPP group configuration is required for threshold calculation at UPF.

UPF uses GTPP group name available from APN configuration. Only one GTPP group should be associated under APN configuration.

The following is a sample configuration:

```
configure
  context context_name
    apn apn_name
      gtpp group group_name
      ip context-name name
      exit
    gtpp group group_name
      gtpp egcdr service-data-flow threshold interval interval
      gtpp egcdr service-data-flow threshold volume downlink bytes
      gtpp egcdr service-data-flow threshold volume uplink bytes
      gtpp egcdr service-data-flow threshold total bytes
      end
```

If any one of the above service-data-flow thresholds is hit for offline URR, the UPF sends SX_SESSION_REPORT_REQUEST towards SMF reporting the data volume.



CHAPTER 18

ECS Regular Expression

- Feature Summary and Revision History, on page 231
- Feature Description, on page 232
- How It Works, on page 232
- Configuring Regex Rule, on page 233
- Sample Configuration, on page 234
- Monitoring and Troubleshooting, on page 234

Feature Summary and Revision History

Summary Data

Table 48: Summary Data

Applicable Product(s) or Functional Area	5G-UPF
Applicable Platform(s)	VPC-SI SMI
Feature Default Setting	Disabled - Configuration Required
Related Changes in This Release:	Not Applicable
Related Documentation	Not Applicable

Revision History

Table 49: Revision History

Revision Details	Release
First introduced.	2022.01.2

Feature Description

The ECS Regular Expression feature supports the implementation of regex engine in the User Plane Function (UPF). Furthermore, this feature allows you to configure the regex rule through RCM.

The UPF supports the following protocols as part of regex engine rebuild and rule matching.

- HTTP
  - URL
  - URI
  - HOST
- WWW
  - URL
  - URI
- RTSP
  - URL
  - URI

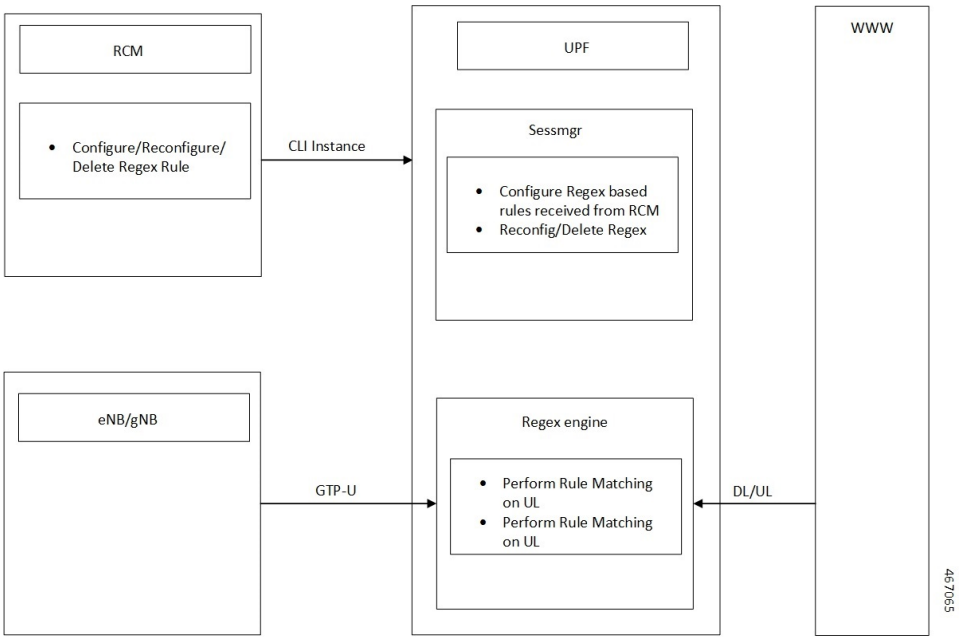
How It Works

The following table lists the special characters that you can use in regex rule expressions.

Convention	Description
*	Zero or more characters.
+	Zero or more repeated instances of the token preceding the +.
?	Zero or one character.
\character	Escaped character.
\?	Match on a question mark (\<ctrl-v>?)
\+	Match on a plus sign
*	Match on an asterisk
\a	Alert (ASCII 7)
\b	Backspace (ASCII 8)
\f	Form-feed (ASCII 12)
\n	New line (ASCII 10)
\r	Carriage return (ASCII 13)

Convention	Description
\t	Tab (ASCII 9)
\v	Vertical tab (ASCII 11)
\0	Null (ASCII 0)
\\	Back slash
Bracketed range [0-9]	Matching any single character from the range.
A leading ^ in a range	No match in the range. All other characters represent themselves.
.\x##	Any ASCII character as specified in two-digit hex notation. For example, \x5A yields a ‘Z’.

The following diagram illustrates the regex rule configuration through RCM:



Configuring Regex Rule

Configure the regex rule through RCM using UPF CLI instance or directly on the UPF through CLI.

```
configure
  require rcm-configmgr
  end
```

Sample Configuration

Following is a sample configuration for configuring the Regex Rule.

```
configure
   active-charging service <service_name>
      ruledef <ruledef_name>
         http url regex <regex_url>
         rtsp uri regex <regex_uri>
         www url regex <regex_url>
         end
```

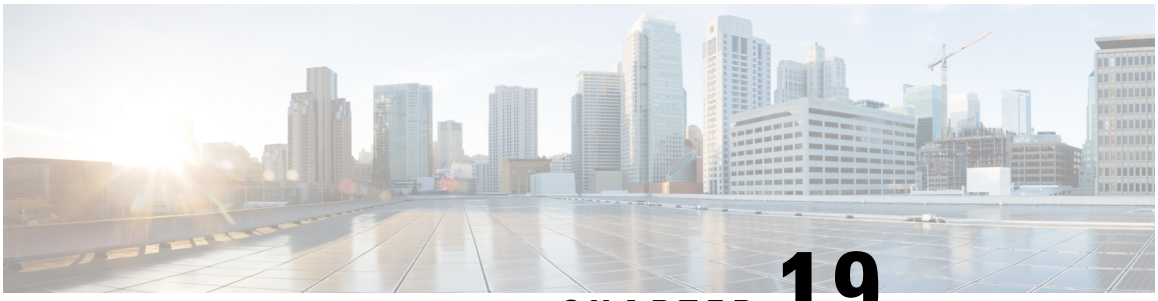
Monitoring and Troubleshooting

This section provides information on CLI commands that are available for monitoring and troubleshooting the feature.

Show Commands and Outputs

This section provides information about show CLI commands that are available in support of ECS Regex feature in UPF.

Show Commands	Description
show user-plane-service regex status	Use this command to display the engine status for the SessMgr instance.
show user-plane-service regex statistics memory	Use this command to display the memory statistics for the SessMgr instance.
show user-plane-service regex statistics memory summary	Use this command to display the combined memory summary for the SessMgr.
show user-plane-service regex statistics ruledef	Use this command to display the regex ruledef statistics for the SessMgr.
show user-plane-service regex statistics ruledef summary	Use this command to display the combined regex ruledef statistics summary for the SessMgr.



CHAPTER 19

GTP-U Support

- Feature Summary and Revision History, on page 235
- Feature Description, on page 236
- How it Works, on page 237
- Enabling Dual Stack Support on S5u/N3 Interface, on page 240
- Disabling UDP Checksum, on page 240

Feature Summary and Revision History

Summary Data

Table 50: Summary Data

Applicable Product(s) or Functional Area	5G-UPF
Applicable Platform(s)	VPC-SI SMI
Feature Default Setting	Enabled – Always-on
Related Changes in this Release	Not Applicable
Related Documentation	Not Applicable

Revision History

Table 51: Revision History

Revision Details	Release
Optimization of UDP checksum is added in this release.	2021.02.0
First introduced	2020.02.0

Feature Description

Table 52: Feature History

Feature Name	Release Information	Description
Dual Stack Support on S5u Interface	2024.02.0	UPF supports dual stack to handle IPv4 and IPv6 connections between SGW-U and PGW-U on the S5u interface. **Default Setting**: Enabled – Always-on
Dual Stack Support on S1-U and N3 Interfaces	2024.01	UPF supports dual stack to handle IPv4 and IPv6 connections on N3 and S1-U interfaces. **Default Setting**: Enabled – Always-on

3GPP specifies provisions for UEs capable of supporting both 5G and 4G NAS to connect to the E-UTRAN and 5G core network.

To forward data (G-PDUs and End Marker packets) during an EPS to 5GS handover, the SMF performs the following tasks:

- Provisions one PDR per E-RAB (that supports data forwarding for at least one QoS flow).
- Creates and associates one QER with each PDR to request the UPF to insert a GTP-U PDU Session Container extension header including the QFI. The PDR includes the QFI IE set to the QFI value of one of the QoS flows mapped to the E-RAB.

Data forwarding during handovers between 5GS and EPS is supported as follows. For reference, see 3GPP TS 38.300.

- For 5G to 4G handover, the source NG-RAN node sends one or several end-markers including one QFI of those QoS flows mapped to the same E-RAB and sends the end-marker packets to the UPF over the PDU session tunnel. UPF removes the QFI and maps to an appropriate E-RAB tunnel toward SGW.
- For 4G to 5G handover, the source eNB forwards the received end markers in the EPS bearer tunnel to the SGW. Then SGW forwards them to the UPF. The UPF adds one QFI among the QoS flows mapped to that E-RAB to the end-markers. Then, the UPF sends those end-markers to the target NG-RAN node in the per PDU session tunnel.

Note UPF supports dual-stack to handle IPv4 and IPv6 connections over the S1-U, S5u, and N3 GTP-U interfaces.

Error Indication and GTP-U Path Failure

The UPF notifies an Error Indication message for a GTP-U peer to the sender when a GTP-PDU is received with a TEID that doesn't exist. This notification ensures that there are no stale sessions or bearers, and maintains consistency in the network.

Error Indication and GTP-U Path Failure between SMF and UPF nodes are supported over the N4 interface. For the neighbor nodes, it's supported over the S1u/S5u interfaces.

Behavior variations of local-purge or signal-peer for Error Indication and GTP-U Path Failure are considered in this implementation.

- When Error Indication is received, the UPF communicates the TEID and GTPU-peer information with the SMF to ensure deletion or modification of the GTPU-peer.
- On receiving a GTP-U packet with nonexisting TEID, the UPF generates and sends Error Indication with TEID and GTP-U peer entries.
- The deletion of a session or a bearer is decided based on the Path Failure detection at SMF or UPF.
- GTP-U Path Failure is detected using GTP-U echo messages between the UPF nodes, and between the UPF and SMF nodes.

How it Works

Call Flows

Initial Attach on E-UTRAN via MME and S-GW

Initial attach on E-UTRAN/EPS follows the procedure defined in 3GPP TS 23.401, Section 5.3.2.1.

The following diagram shows the call flow derived from 3GPP reference for initial attach on E-UTRAN/EPS.

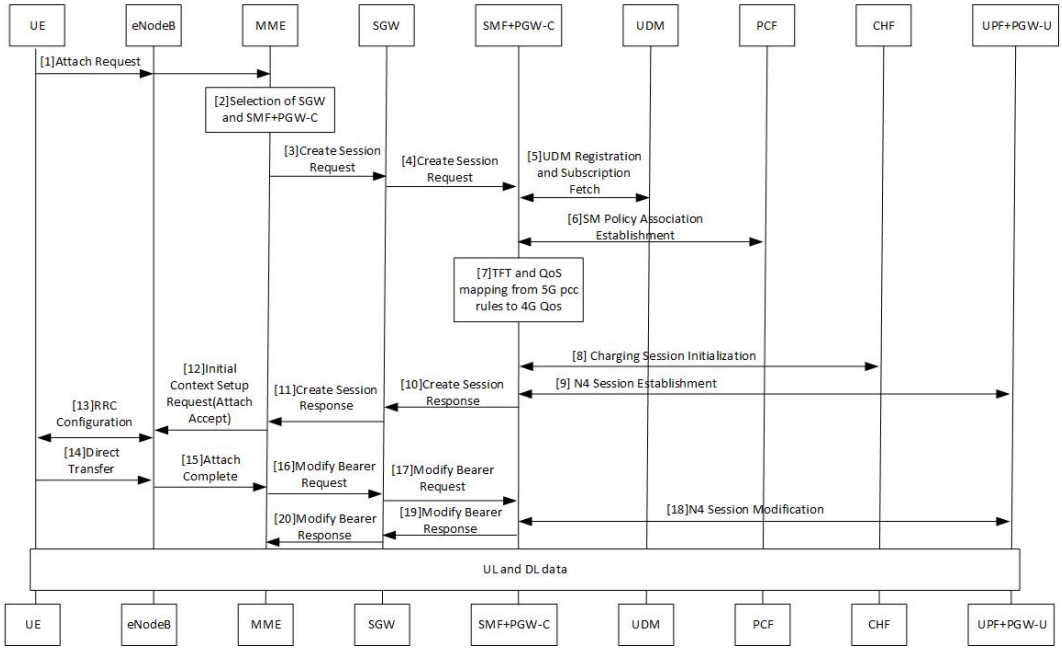


Table 53: Initial Attach on E-UTRAN via 5G Core Call Flow

Step	Description
1	At Step 9, SMF+PGW-C perform a UPF selection and perform N4 Session Establishment procedure. Since this session is a 4G session connecting to SMF+PGW-C, separate CN tunnel is created for each bearer and QFI is not sent in the QER and PDR, correlation ID might be present.
2	At Step 18, SMF+PGW-C performs N4 Session Modification to update the eNodeB TEID on the data path to the UPF.

The 3GPP specifications provide mechanisms to achieve mobility of a UE from LTE to 5G NR and vice versa. This mobility is achieved in two different architectures – with and without N26 interface between AMF and MME.

5G to EPS Handover with N26 Interface

5G to EPS handover with N26 interface is defined in 3GPP TS 23.502, Section 4.11.1.2.1. The following diagram shows the detailed call flow for N26 interface.

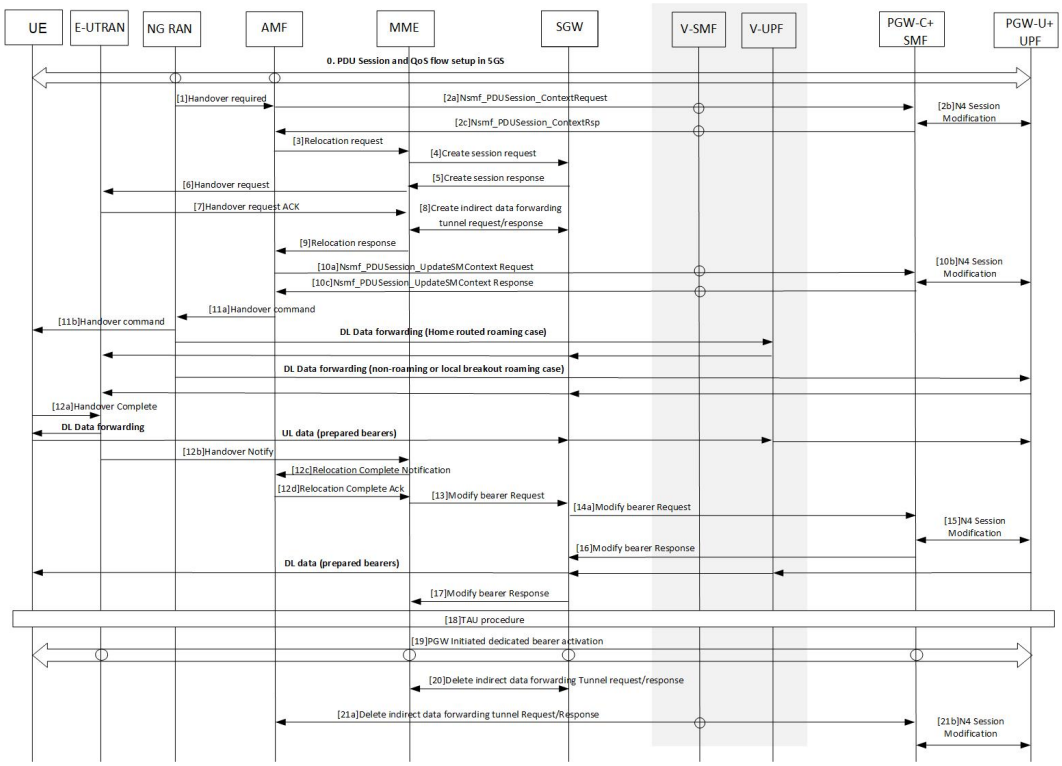


Table 54: 5G to EPS Handover with N26 Interface Call Flow

Step	Description
1	In Step 2b, the SMF+PGW-C sends the N4 Session modification to the UPF to establish the CN tunnel for each EPS bearer. The bearer mapping to the 5G QoS and PCC rules received from PCC is already present with SMF. The SMF also contains the bearer IDs obtained from the Bearer ID Allocation procedure. SMF+PGW-C creates new PDRs for the N4 session and gets TEID allocated for each bearer as required by the 4G system.
2	In Step 10b, SMF+PGW-C sends N4 Modification Request to UPF to create additional PDRs and FARs to receive the redirected DL data over the indirect tunnel from NG RAN and forward them to eNodeB. The uplink PDRs in this case has the QFI to match forwarded DL data from NG RAN and the associated QER does not have the QFI as data needs to be forwarded to eNodeB. Also, the FAR redirects the received data to eNodeB over appropriate tunnel based on the QFI.
3	At Step 11, for the QoS flows indicated in QoS Flows for Data Forwarding, NG-RAN initiates data forwarding through the UPF based on the CN Tunnel Info for Data Forwarding per PDU Session. Then the UPF maps data received from the data forwarding tunnel(s) in the 5GS to the data forwarding tunnel(s) in EPS and sends the data to the target eNodeB through the Serving GW.
4	In Step 15, the SMF sends N4 Modification Request to UPF to activate the DL data path to E-UTRAN. At this time, both the indirect tunnel and the direct DL path are activated towards eNodeB.
5	At Step 21, the SMF sends N4 Modification Request to the UPF to delete the indirect forwarding tunnel.

Other call flows related to EPS to 5G and 5G to EPS handover with N26 interface, or without N26 interface are defined in 3GPP 23.502, Section 4.11.1.2.1 and Section 4.11.2.

Error Indication Handling on UPF

UPF, on receiving Error Indication, initiates a PFCP Session Report Request with Error Indication Report that includes remote F-TEID containing TEID and GTP-U Peer address.

- For PGW-U, Error Indication message is sent or received over S5u.
- For SAEGW-U, Error Indication message is sent or received over S1u.
- For SGW-U, Error Indication message is sent and received over S1u and S5u.

UPF generates Error Indication with TEID and GTP-U Peer Address towards a peer when a data packet is received with TEID for which a session or bearer doesn't exist.

GTP-U Path Failure Support at UPF

GTP-U Echo Requests is initiated and sent periodically as per the configured interval on UPF. GTP-U Echo Response is sent for the GTP-U Echo Request received from SMF over GTP-U tunnel.

If Response is not received for the GTP-U Echo Request, the UPF retries Echo Requests based on configured retransmission timeout and maximum retries. When retries are exhausted, the UPF initiates PFCP node

Report Request including (Node ID, Node Report Type, User Plane Path Failure Report including Remote GTP-U Peer).

If UPF receives PFCP Node Report Response and PFCP Session Deletion Request to delete the session, it responds to the deletion request with usage reports.

Enabling Dual Stack Support on S5u/N3 Interface

The following is a sample configuration to enable dual stack on S5u/N3 interface:

```
context SAEGW
    interface pgw-gtpu loopback
      ip address 81.81.20.5 255.255.255.255
    exit

    gtpu-service pgw-gtpu
      no udp-checksum
      bind ipv4-address 81.81.21.x ipv6-address 2001:4888:81:81::21:5 bearer-type all
    exit

    user-plane-service user_plane_svc
      associate gtpu-service pgw-gtpu upf-ingress
      associate gtpu-service SxC cp-tunnel
      associate sx-service sx-svc
      associate fast-path service
      associate control-plane-group SAEGW
    exit
```

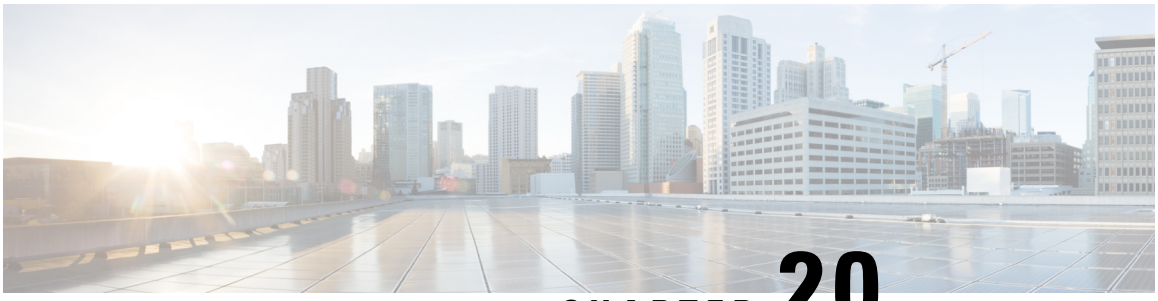
Disabling UDP Checksum

This functionality disables the UDP checksum in UDP header of the GTP-U packet. The value of the UDP checksum is set to zero.

Disabling UDP Checksum

Use the following configuration to disable the UDP checksum in UDP header of the GTP-U packet.

```
configure
  context context_name
    gtpu-service service_name
      no udp-checksum
      end
```



CHAPTER 20

Heartbeat Support for N4/Sx Interface

- Feature Summary and Revision History, on page 241
- Feature Description, on page 242
- How It Works, on page 242
- Configuring Heartbeat for N4/Sx Interface, on page 243
- Monitoring and Troubleshooting, on page 244

Feature Summary and Revision History

Summary Data

Table 55: Summary Data

Applicable Product (s) or Functional Area	5G-UPF
Applicable Platforms	VPC-SI SMI
Feature Default Setting	Disabled – Configuration Required
Related Changes in this Release	Not Applicable
Related Documentation	*UCC 5G UPF Configuration and Administration Guide*

Revision History

Revision Details	Release
First introduced.	2021.01.0

Feature Description

In accordance with 3GPP TS 29.244, support has been added for node-level Heartbeat procedures between the Session Management Function (SMF) and User Plane Function (UPF) over N4/Sx interface.

The Heartbeat procedure contains the following two messages:

1. Heartbeat Request
2. Heartbeat Response

Heartbeat Request

The SMF or the UPF sends a Heartbeat Request on a path to the peer node to find out if it is alive. The Heartbeat Request messages are sent for each peer with which a Packet Forwarding Control Protocol (PFCP) control association is established.

For each peer with which a PFCP control association is established, an SMF or UPF is prepared to receive a Heartbeat Request at any time, and replies with a Heartbeat Response.

Heartbeat Response

This message is sent as a response to a Heartbeat Request.

How It Works

The SMF and UPF sends Heartbeat messages after configurable time duration. If the peer does not respond, the message is retried for configured number of times with the retry-interval and then the configured action is taken for the calls associated with the corresponding peer.

Recovery Time Stamp Information Element (IE), which contains the start time of the node, is supported by both Heartbeat Request and Heartbeat Response. Heartbeat Request contains its own Recovery Time Stamp value and sends it to the peer while Heartbeat Response contains the peers Recovery Time Stamp value.

Path Failure Detection

Path failure is detected in following conditions:

1. Heartbeat failure: This condition occurs when the peer does not respond to the Heartbeat that is sent and also retires.
2. Recovery Time stamp change in Heartbeat: This condition occurs when the Heartbeat Request or Heartbeat Response has a new larger value than the previously received value.
3. Recovery Time stamp change in N4/Sx Association message: This condition occurs when the N4/Sx association message is received again from the peer with a new Recovery Time Stamp.

Path Failure Handling

When the Recovery Time Stamp value received is more than the previously received value, then the peer restart is detected. If the Recovery Time Stamp value is lower than the previously received value then the value is ignored and peer restart is not detected.

When a peer restart is detected, an SNMP Trap is generated to indicate the path failure for the peer. Also, based on the path failure configuration (refer Configuring Heartbeat for N4/Sx Interface, on page 243), all the calls connected to that peer can be cleared.

Configuring Heartbeat for N4/Sx Interface

This section provides information about the CLI commands available in support of this feature.

Enabling Heartbeat for Sx Interface

Use the following commands under Sx Service Configuration mode to enable Heartbeat parameters for N4/Sx interface.

```
configure
   context context_name
      sx-service service_name
         [ default ] sx-protocol heartbeat { interval seconds |
max-retransmissions number | path-failure detection-policy {
control-recovery-timestamp-change | heartbeat-retry-failure |
heartbeat-recovery-timestamp-change } | retransmission-timeout seconds }
       no sx-protocol heartbeat { interval | path-failure detection-policy
 { control-recovery-timestamp-change | heartbeat-retry-failure |
heartbeat-recovery-timestamp-change }
        end
```

Notes:

- **default**: Sets/restores default value assigned for specified parameter.
- **no**: Disables the followed option.
- **heartbeat**: Configures N4/Sx Heartbeat parameters.
- **interval** *seconds*: Configures Heartbeat interval (in seconds) for N4/Sx Service. *seconds* must be an integer in the range of 1 to 3600.
- **max-retransmissions** *number*: Configures maximum retries for N4/Sx Heartbeat request. Must be followed by integer, ranging from 0 to 15. Default is 4.
- **retransmission-timeout** *seconds*: Configures the Heartbeat retransmission timeout for N4/Sx service, in seconds, ranging from 1 to 20. Default is 5.
- **path-failure**: Specifies the policy to be used when path failure happens through Heartbeat request timeout.

Configuring Detection Policy for Path Failure

Use the following commands under Sx Service Configuration mode to specify detection policy to be used for path failure.

```
configure
   context context_name
      sx-service service_name
       [ default | no ] sx-protocol heartbeat path-failure detection-policy
 { control-recovery-time-stamp-change | heartbeat-retry-failure |
heartbeat-recovery-
timestamp-change }
        end
```

NOTES:

- **default**: Sets/restores default value assigned for specified parameter.
- **no**: Disables the followed option.
- **detection-policy**: Specifies the policy to be used. Default action is to do cleanup upon Heartbeat request timeout.
- **control-recovery-time-stamp-change**: Path failure is detected when the recovery timestamp in control request/response message changes.
- **heartbeat-retry-failure**: Path failure is detected when the retries of Heartbeat messages times out.
- **heartbeat-recovery-timestamp-change**: Path failure is detected when the recovery timestamp in Heartbeat request/response message changes.

Monitoring and Troubleshooting

This section provides information about CLI commands available for monitoring and troubleshooting the feature.

Show Command(s) and/or Outputs

This section provides information regarding show commands and/or their outputs in support of this feature.

show sx-service all

The output of this show command has been enhanced to include the following fields introduced in support of this feature.

- SX Heartbeat
  - Interval
  - Retransmission Timeout
  - Max Retransmission
- SX path failure detection policy

- Heartbeat Timeout
- Heartbeat Req/Rsp Recovery timestamp change
- Control Msg Recovery timestamp counter change

show sx-service statistics all

The output of this show command has been enhanced to include the following fields introduced in support of this feature.

- Heartbeat Request
  - Total TX
  - Total RX
  - Initial TX
  - Initial RX
  - Retrans TX
- Heartbeat Response
  - Total TX
  - Total RX

Disconnect Reasons

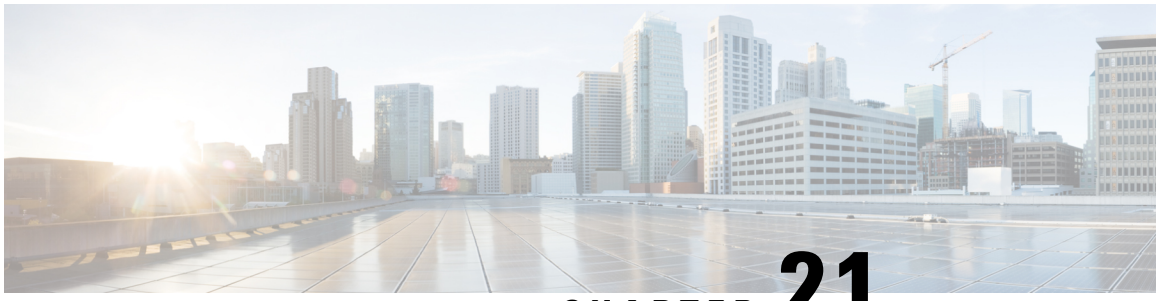
The following disconnect reason has been added in support of this feature:

- sx-path-failure - When the Recovery timestamp changes or heartbeat failure is detected, based on the configuration, calls are cleared with this disconnect reason.

SNMP Traps

The following SNMP traps have been added in support of this feature:

- SxPathFailure - This trap is generated when the peer path failure is detected.
- SxPathFailureClear - This trap is generated when the path is restored for the peer.



CHAPTER 21

Idle Mode Buffering and Paging

This chapter covers the following topics:

- Feature Summary and Revision History, on page 247
- Feature Description, on page 248
- Buffering Action Rule Call Flow, on page 248
- Downlink Data Report for First DL Packet, on page 249
- Paging Policy Differentiation, on page 249

Feature Summary and Revision History

Summary Data

Table 56: Summary Data

Applicable Product(s) or Functional Area	5G-UPF
Applicable Platform(s)	VPC-SI SMI
Feature Default Setting	Enabled – Always-on
Related Changes in this Release	Not Applicable
Related Documentation	Not Applicable

Revision History

Table 57: Revision History

Revision Details	Release
First introduced.	2020.02.0

Feature Description

A Buffering Action Rule (BAR) provides instructions to control the buffering behavior of the User Plane Function (UPF). The BAR controls the buffering behavior for all Forwarding Action Rules (FARs) of the Packet Forwarding Control Protocol (PFCP) session. This control is applicable when the PFCP session is set with an Apply Action parameter, which requests packets to be buffered and associated with the respective BAR.

How it Works

If the User Plane Function indicates the support of the feature UL or DL Buffering Control (UDBC), the SMF provides the buffering packet count IE in a BAR. The buffering count IE is created during a PFCP Session Establishment procedure or a PFCP Session Modification procedure. The SMF modifies it in a subsequent PFCP session modification request, "and" or "or" a PFCP Session Report Response message. The same BAR associates with all the FARs in a PFCP session to indicate that all service data flows in the PFCP session shares the same buffer in the UPF for the PFCP session. One BAR is created per PFCP session.

Provisioning of Buffering Action Rule in the UPF

The SMF provisions multiple buffering parameters in a BAR. It is in Create BAR or Update BAR in various PFCP messages.

Currently, UPF supports the following IE:

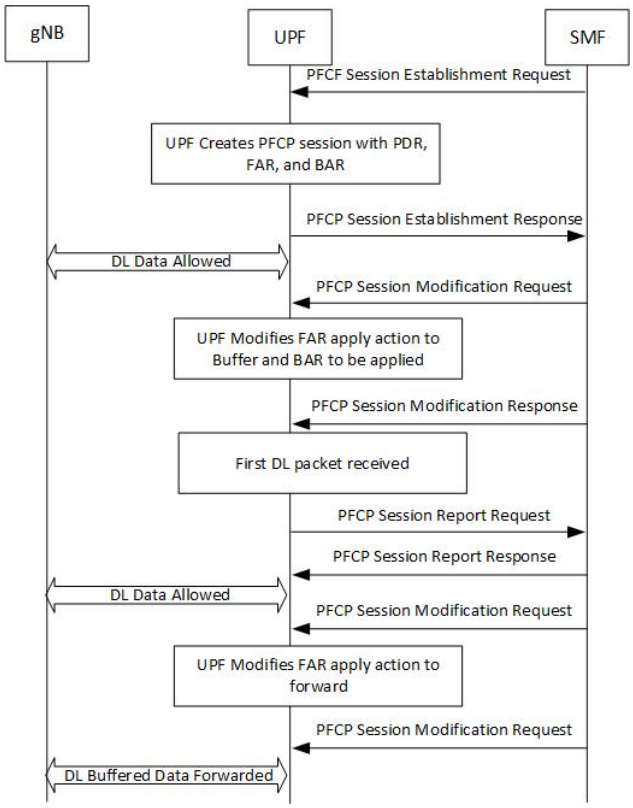
- The suggested buffering packet count IE—If the UPF indicates the support of the feature UDBC to indicate the number of packets. It includes both uplink or downlink that the SMF suggests buffering in the UPF, until it receives new instructions from the SMF. Example: when the new quota is granted.
- DL buffering suggested packet count IE—This IE is received with update BAR from SMF in Session Report Response message, if SMF wants more DL packets to be buffered on UPF.

The UPF does not apply the DL buffering duration and DL buffering suggested packet count parameters and deletes these parameters from the BAR (without explicit request from the SMF) when extended buffering of downlink data packets ends in the UPF. The UPF does not apply buffering when it receives the new instruction from the SMF. The buffered packets are either dropped or forwarded following the packet forwarding model and considering that the buffered packets are already processed earlier.

Buffering Action Rule Call Flow

This section describes the provisioning of buffering action rule in the UPF call flow.

Figure 12: Buffering Action Rule



Downlink Data Report for First DL Packet

When instructed to buffer and notify the SMF about the arrival of a DL packet, the UPF notifies the SMF, when it receives a first downlink packet for a given FAR. The UPF notifies the DL packet arrival by sending a PFCP Session Report Request including a Downlink Data Report IE identifying the PDR(s) for which downlink packets was received.

Paging Policy Differentiation

The UPF supports the Paging Policy Differentiation, for each PDR and for each packet that triggers a Downlink Data Notification, the UPF function copies the value of the DSCP in ToS (IPv4) or TC (IPv6) information received in the IP payload in Downlink Data Service Information IE.

For each PDR and for each packet that triggers a Downlink Data Notification, if the QFI of the downlink data packet is available, the QFI is also sent in Downlink Data Service Information IE.

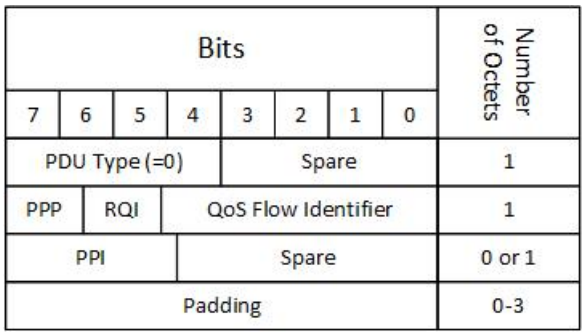
Paging Policy Indicator (PPI)

The SMF sends the PPI value in Create QER or Update QER, if UPF needs to set Paging Policy Indicator in outgoing PDU packets.

Frame Format for the PDU Session User Plane Protocol

Downlink PDU Session Information (PDU Type 0) - This frame format is defined to allow the NG-RAN to receive some control information elements which are associated with the transfer of a packet over the interface. The following figure shows the respective DL PDU SESSION INFORMATION frame.

Figure 13: DL PDU SESSION INFORMATION (PDU Type 0) Format



QoS Flow Identifier (QFI)

Description: When present, the QoS Flow Identifier (QFI) parameter indicates the QoS Flow Identifier of the QoS flow to which the transferred packet belongs.

Value Range: The value range is between 0 to 2 6 -1.

Field Length: 6 bits.

Paging Policy Presence

Description: The Paging Policy Presence (PPP) parameter indicates the presence of the Paging Policy Indicator (PPI).

Value Range: A value of 0 indicates that Paging Policy Indicator is not present and 1indicates that Paging Policy Indicator is present.

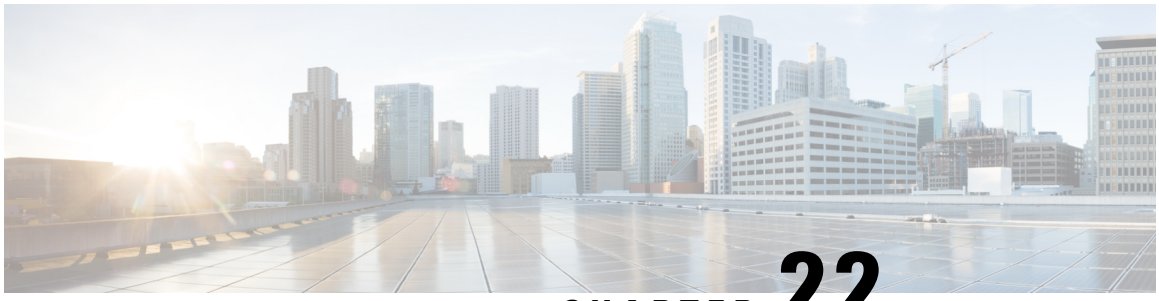
Field Length: 1 bit.

Paging Policy Indicator

Description: When present, the Paging Policy Indicator (PPI) is used for paging policy differentiation (see details in 3GPP TS 23.501). This field applies to PDU sessions of IP type.

Value Range: the value range is between 0 to 2 3 -1.

Field Length: 3 bits.



CHAPTER 22

Indirect Forwarding Tunnel

- Revision History, on page 251
- Feature Description, on page 251
- How It Works, on page 251
- Configuring Indirect Forwarding Tunnel, on page 254
- Monitoring and Troubleshooting, on page 255

Revision History

Revision Details	Release
First introduced	2021.02.0

Feature Description

The UPF supports Indirect Forwarding Tunnel (IDFT) procedures for creation and deletion, which are applicable for Pure-S and Collapsed calls with dedicated bearers. This feature is applicable for IDFT support with S-GW Relocation.

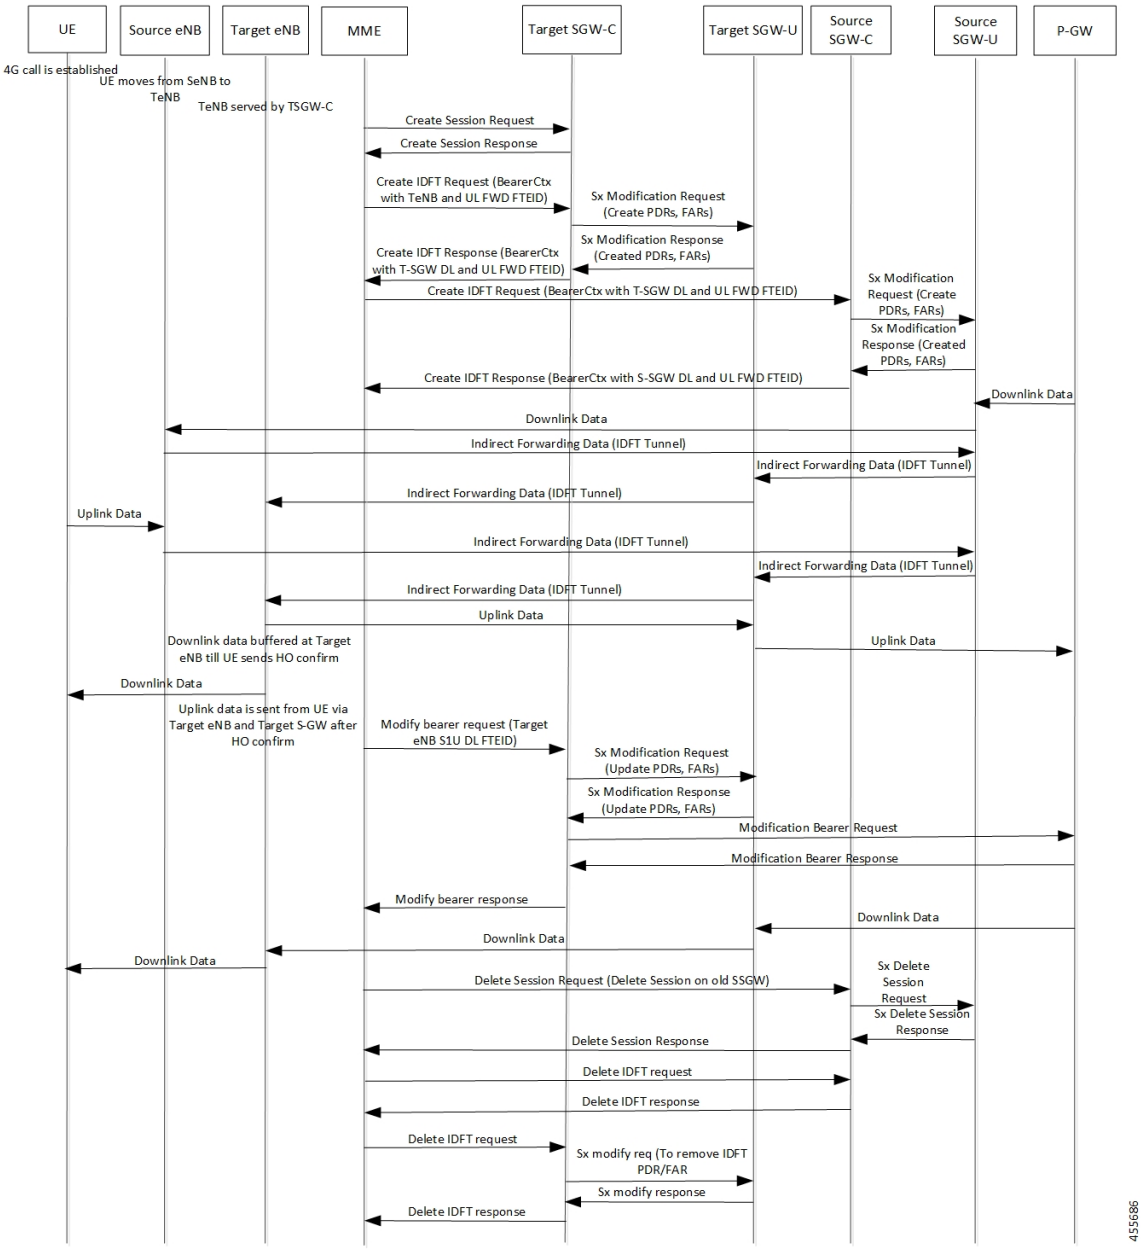
Note The IDFT in UPF is a CLI-controlled feature. By default, the IDFT feature in UPF is disabled.

How It Works

Call Flow

The following call flow illustrates the IDFT support with S-GW Relocation.

Figure 14: IDFT Support with S-GW Relocation



The above call flow describes the IDFT tunnels establishment and deletion with S-GW relocation and without MME change.

If IDFT tunnels are not deleted by MME, then S-GW initiates the local delete of IDFT tunnels.

This feature supports the following scenarios for the Pure-S and Collapse calls:

- S-GW relocation with same MME
- S-GW relocation with same MME and different eNodeB
- S-GW relocation with different MME

- S1-based eNodeB Handoff
- EUTRAN to UTRAN Handoff
- EUTRAN to UTRAN Handoff with S-GW relocation
- UTRAN to EUTRAN Handoff
- UTRAN to EUTRAN Handoff with S-GW relocation
- Tracking Area Update (TAU) with S-GW change and indirect data forwarding
- Radio Access Bearer (RAB) during Active IDFT
- Sx transaction timeout during IDFT setup or removal
- Pending Sx transaction (event from PCRF or OCS) and IDFT request comes in
- Create Bearer Request (CBR) during Active IDFT
- Update Bearer Request (UBR) during Active IDFT
- Delete Bearer Request (DBR) during Active IDFT
- Modify Bearer Request (MBR) behavior on other PDN during Active IDFT
- Source MME path failure
- Target MME path failure
- MME path failure with NTSR enabled
- eGTP-C S5 path failure
- eGTP-C S5 path failure with P-GW restart notification enabled
- Sx path failure (clean IDFT and calls)
- Abort session (clear sub all, local abort, and so on.)
- CBR, UBR on other PDN during Active IDFT
- DBR on other PDN/bearer during Active IDFT
- S1-u path failure for target eNodeB
- S1-u path failure for source eNodeB
- S-GW path failure for source S-GW
- S-GW path failure for target S-GW
- S1-u error indication on the default bearer while Active IDFT
- S1-u error indication on the dedicated bearer while Active IDFT
- S1-u error indication from the target S-GW to source S-GW bearer
- S1-u error indication from the target eNodeB to target S-GW bearer
- Sending End Marker when tearing down IDFT tunnel after failure
- If SMF ICSR/SR leads to the cleanup of IDFTs, the UPF also cleans the IDFT PDRs/FARs

5G to 4G Handover with IDFT

In compliance with 3GPP TS 23.502 v15.5.1, the 5G to 4G handover with IDFT is supported in UPF. Refer the "5GS to EPS handover using N26 interface" section in the 3GPP specification for details about call/datapath flow for IDFT.

This functionality isn't CLI-controlled in UPF.

4G to 5G Handover with IDFT

In compliance with 3GPP TS 23.502 v15.5.1, the 4G to 5G handover with IDFT is supported in UPF. Refer the "EPS to 5GS handover using N26 interface" (preparation and execution phase) section in the 3GPP specification for details about the call flow.

This functionality isn't CLI-controlled in UPF.

Supported Functionality

The IDFT feature supports the following functionality:

- Create IDFT request for Collapsed, Pure-S, combination of Collapsed and Pure-S multi-PDN calls with multiple bearers.
- Data transfer on downlink and uplink IDFT bearers.
- Deletion of IDFT request from MME. Also, timer-based deletion of IDFT bearer after expiration of a default value of 100 seconds, if the MME does not send an IDFT request for deletion.
- Deletion of IDFT PDN, including Clear/Delete subscribers from MME/P-GW, when normal PDN goes down.
- IDFT creation of Sx Failure Handling for Pure-S and Collapsed PDN.



Important Transport GTP-U address capability is assumed to be same across eNodeB and S-GW.

Limitations

The IDFT feature has the following limitations:

- Message interaction and collision during IDFT PDN establishment or deletion with any other procedure is not supported.
- S11/S5 and Sx Path Failure Handling on non-IDFT PDN is not supported when IDFT PDN is Active.
- Deletion of partial dedicated bearers in IDFT connected-state is not supported.

Configuring Indirect Forwarding Tunnel

This section describes the CLI commands available in support of IDFT feature.

Enabling Indirect Forwarding Tunnel Feature

On SMF, use the following CLI commands to enable or disable the IDFT feature.

```
configure
  context context_name
    sgw-service service_name
      [ default | no ] egtp idft-support
      end
```

NOTES:

- **idft-support**: Enables/Disables the IDFT feature in UPF.
- By default, the IDFT feature is disabled and this CLI command is applicable on run-time change.

Verifying the Indirect Forwarding Tunnel Feature

show sgw-service name <service_name>

On SMF, the output of this CLI command has been enhanced to display if the IDFT feature is enabled or disabled.

- IDFT-Feature Support for CUPS: Enabled/Disabled

Monitoring and Troubleshooting

This section provides information regarding the CLI commands available in support of monitoring and troubleshooting the feature.

Show Commands Input and/or Outputs

This section provides information regarding show commands and their outputs in support of the feature.

show subscribers saegw-only full all

On UPF, use this command to see the IDFT Local and Remote TEID data. The following is a sample output:

```
Indirect Fwding   : Active
DL fwd local  addr: 209.165.201.28          DL fwd remote   addr: 209.165.201.27
DL fwd local  teid: [0x80028004]            DL fwd remote   teid: [0x2002d2e5]
UL fwd local  addr: 209.165.201.28          UL fwd remote   addr: 209.165.201.27
UL fwd local  teid: [0x8002a004]            UL fwd remote   teid: [0x20042bca]
```

show subscribers user-plane-only callid <call_id> pdr all

On UPF, use this command to see the PDRs created for IDFT. The following is a sample output:



Important IDFT PDRs will have ACCESS as the source and destination interface type.

```
+-----Source Interface:          (C) - Core            (A) - Access
|-----Type                       (P) - CP-function     (.) - Unknown
|
|+----Destination Interface:     (C) - Core            (A) - Access
||----Type                       (P) - CP-function     (.) - Unknown
||
||
||+---Rule-Type:                 (S) - Static          (P) - Predefined
|||---Type                       (D) - Dynamic         (.) - Unknown
|||
|||
vvv   PDR-ID    Associated FAR-ID    Associated URR-ID(s)    Associated QER-ID(s)
---   ------    -----------------    --------------------    --------------------
ACS   0x0001         0x8001              n/a                  0x80000001
CAS   0x0002         0x8002              n/a                  0x80000001
ACD   0x0003         0x0003           0x00000007              0x00000002
                                         n/a                  0x80000003
CAD   0x0004         0x0004           0x00000007              0x00000002
                                         n/a                  0x80000003
CAD   0x0005         0x0005           0x00000000              n/a
ACD   0x0006         0x0006           0x00000000              n/a
CAD   0x0007         0x0007           0x00000000              n/a
ACD   0x0008         0x0008           0x00000000              n/a
AAD   0x0009         0x0009           0x00000000              n/a
AAD   0x000A         0x000A           0x00000000              n/a
AAD   0x000B         0x000B           0x00000000              n/a
AAD   0x000C         0x000C           0x00000000              n/a

Total subscribers matching specified criteria: 1
```

Similarly, you can use the **show subscribers user-plane-only callid** *call_id* **far all** CLI command to see the FARs created for IDFT

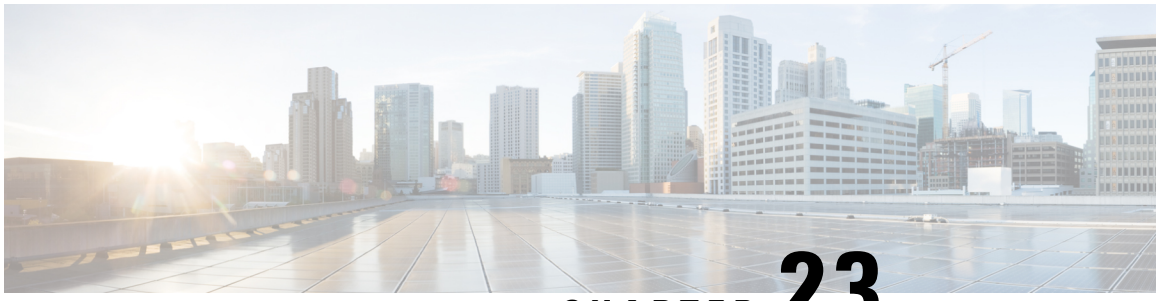
show subscribers user-plane-only full all

Important Data statistics on IDFT PDRs are captured in the same way as existing PDR statistics. However, it is captured with a limitation – Statistics for DL and UL IDFT will be incremented in Pkts-Down and Bytes-Down category.

The following is sample output:

```
Static & Predef Rule Match stats:
Rule Name            Pkts-Down  Bytes-Down Pkts-Up    Bytes-Up   Hits       Match-Bypassed
  FP-Down(Pkts/Bytes)  FP-Up(Pkts/Bytes)
-------------------- ---------- ---------- ---------- ---------- ---------- --------------
  -------------------  -----------------
catchall                      0          0          3       1368          3              0
          0/0              0/0

Dynamic Rule Match stats:
PDR Id   Pkts-Down  Bytes-Down Pkts-Up    Bytes-Up   Hits       Match-Bypassed
FP-Down(Pkts/Bytes)    FP-Up(Pkts/Bytes)
-------- ---------- ---------- ---------- ---------- ---------- --------------
-------------------    -----------------
0x0004        2       856          0          0          2              0              0/0
                 0/0
0x000b        2       856          0          0          2              0              0/0
                 0/0
0x000c        2       168          0          0          2              0              0/0
                 0/0
```



CHAPTER 23

IP Allocation through DHCP Server

Table 58: Feature History

Feature Name	Release Information	Feature Description
DHCP-based IPv6 address allocation without prefix limitations	2025.03.0	This feature allows the UPF to support IPv6 address allocation to the subscriber UEs through external DHCP server, without prefix limitations. **Commands introduced:** **disable ipv6-validation**: This CLI is configured under DHCP Client Profile configuration mode to disable IPv6 validation. **Commands enhanced:** **disable dhcpv6-client-unicast slot/port** *slot_port_number*: This CLI is configured under DHCP Client Profile configuration mode to disable unicast option and thus enable multicast. **Default Settings:** Disabled—Configuration Required to Enable
DHCP-based IP Address Allocation	2025.02.0	This feature allows the network operator to use an external DHCP server to receive IP addresses for subscriber UEs. Allocating IP addresses using a DHCP server enables the enterprises with complex 5G networks to manage and control the IP address allocation centrally. **Commands Introduced:** **dhcp server enable-discover slot/port** *slot_port_number* **dhcp-node-id** *node_id*: This CLI is configured under DHCP service configuration mode to allow UPF to discover the DHCP server. **Commands Enhanced:** **dhcp ip skip-validation**: This CLI is configured under the DHCP service configuration mode to skip the IPv4 address validation. **Default Settings:** Disabled—Configuration Required to Enable

In a 5G network, SMF handles the dynamic IP address allocation from the local IPAM pool. However, the complex 5G enterprise networks require centralization of IP address allocation and management.

To facilitate this, UPF supports using a DHCP server to allocate the IP addresses to the subscriber UEs. It allows the network operator to have a centralized IP management that aligns with the needs of complex and large-scale enterprise environments.

- How UPF allocates the UE IP through DHCP server, on page 258
- Enabling DHCP-based IP allocation, on page 263
- Limitations, on page 267

How UPF allocates the UE IP through DHCP server

These stages show the process in which the UPF allocates the UE IP address through external DHCP server:

1. **Session establishment request:** The SMF sends the UE IP Address IE in PDI IE within the PFCP Session Establishment Request message to the UPF, if the SMF is configured to obtain UE IP addresses from the UPF.

   This message contains the VLAN ID(s) in the Create PDR IE in UE IP Address Pool Identity IE and either or both of these message bits are set to one:

   - **CHOOSE IPV4 (CHV4):** SMF requests the UPF for IPv4 address allocation.
   - **CHOOSE IPV6 (CHV6):** SMF requests the UPF for IPv6 address allocation.

2. **DHCP subnet selection:** UPF uses the received VLAN ID(s) for selecting the subnet for UE IP based on some conditions.

   Learn more about the conditions in the DHCP subnet selection criteria topic.

3. **DHCP configuration validation:** UPF validates if the DHCP service is configured and running by checking the configuration validation criteria.

   To know more about the DHCP configuration, see the Enabling DHCP-based IP allocation section.

   If the DHCP configuration validation is successful, UPF accepts the PFCP Establishment Request and initiates a DHCP Discovery handshake toward DHCP server in the given subnet using the received VLAN ID. If the configuration validation fails based on the criteria, the UPF rejects the Establishment Request.

   To know more about the configuration validation, see the Criteria for DHCP configuration validation section.

4. **DORA or SARR handshake process initiation:** UPF initiates the Discovery-Offer-Request-Acknowledge (DORA) or Solicit-Advertise-Request-Reply (SARR) handshake process for IPv4 or IPv6 address allocation respectively.

   UPF broadcasts a **Discovery** message in case of DHCPv4 and multicasts a **Solicit** message in case of DHCPv6 in the subnet. UPF acts as the DHCP client Relay and uses the received VLAN information for sending the Discovery or Solicit messages.

Note

- The UPF receives the unicast response from the DHCP server on the destination port 67.
- UPF acts as client when it sends the broadcast DHCPDISCOVER message in the VLANs source port—68 and the dst port—67.
- For IPv6 prefix allocation, all the messages from UPF to the DHCPv6 server include the IA_PD option, with a requested prefix length of 64.

5. **Offer or Advertise messages:** In case of IPv4, UPF receives the **Offer** message from the DHCPv4 server with an IPv4 address, lease time, and the server identifier.

   In case of IPv6, the UPF receives a unicast **Advertise** message from the DHCPv6 server along with client IPv6 address, lease time, and server identifer.

6. **IPv4 and IPv6 validation:** UPF validates the IP addresses received from the DHCP server. The validation process differs based on the requested type of IP address, that is, IPv4 or IPv6.

   - **IPv4 validation:** If the CLI **[no] skip-validation** is configured, UPF validates the received IPv4 address with the SMF allocated IP chunks.

     If the IPv4 validation is successful, UPF sends a Request message to the DHCP server. If the IPv4 validation fails, UPF does not send the Request message to the DHCPv4 server and the IPv4 address allocation process fails.

     For an IPv4v6 Session, if IPv4 validation fails, UPF sends Sx Session Establishment Response with cause code **Partial Success** with IPv6 address allocated to the UE.

   - **IPv6 validation:** If the CLI **enable ipv6-validation** is configured, UPF sends a **Request** message to the DHCPv6 server. The DHCPv6 server sends a **Reply** message to the UPF.

     For IPv6, prefix validation is done after receiving the Reply message in the SARR handshake.

     If the IPv6 validation is unsuccessful, UPF will send a DHCP DECLINE message to DHCPv6 server to indicate that the UPF will not use the IPv6 address given by DHCPv6 server.

     For an IPv4v6 Session, if IPv6 validation fails, UPF sends Sx Session Establishment Response with cause code **Partial Success** with IPv4 address allocated to the UE.

   Know more about the criteria and process of IPv4 and IPv6 validation addresses in the Validation of IP addresses from DHCP server and SMF section.

7. **Renew and rebind:**

   - **Renew request:** UPF tries to renew the lease when the allocated lease time is about to expire. In this case, UPF unicasts a renew request to the DHCP server at 50% (T1 threshold) of the allocated lease time.

     If the DHCP server does not respond to the renew request, the UPF retries the renew request until the maximum retries exhaust.

     If the server replies with NAK (Negative Acknowledgement), it implies that the lease cannot be extended. The UPF then purges the call and informes the SMF.

     For the IPv6 address, the renew request goes at the T1 timer value in the Reply packet.

- **Rebind request**: If the DHCP server does not respond to the renew request even after the maximum retries exhaust, the DHCP server broadcasts the rebind request when the allocated lease is at 88% (T2 threshold) of the allocated lease time.

  If the DHCP server does not respond to the rebind request, the UPF retries the rebind request until the maximum retries are exhausted. If the DHCP server does not respond even after maximum retries are exhausted, and the allocated lease gets expired, the call gets purged. The UPF notifies about the call deletion to the SMF.

  If the server replies with NAK, it implies that the lease cannot be extended. The UPF then purges the call and informs the SMF.

  For IPv6 address, the rebind packet goes at the T2 timer values in the Reply packet.

The values of the T1 and T2 thresholds are configurable.

To know more about configuring T1 and T2 thresholds, see the *DHCP Service Configuration Mode Commands* chapter in the *Command Line Interface Reference, Modes C - D, StarOS Release 21.28* guide.

Note EPFAR configuration is mandatory to handle renew/rebind failure, and if not configured UPF will do a local purge leading to mismatch between SMF and UPF.

8. **Session deletion:** UPF locally purges the session and sends an indicationtion to SMF that UPF has deleted the call in session report request, once the lease time is expired.

   SMF sends a PFCP session Deletion request over the N4 interface to UPF. Upon receiving the PFCP Session Deletion Request message from SMF, the UPF initiates a fire-and-forget DHCP release request toward the DHCP server.

DHCP subnet selection criteria

The UPF interprets the UE IP Address Pool Identity IE as the VLAN ID for selecting IPv4 subnet according to these conditions:

When N4 session establishment contains...	And Create PDR IE contains ...	Then ...
only CHOOSE IPV4 bit	single UE IP Address Pool Identity IE	UPF interprets as VLAN to identify the subnet for IPv4 IP address allocation.
only CHOOSE IPV4 bit	more than one UE IP Address Pool Identity IEs	UPF ignores additional occurrences of the UE IP Address Pool Identity IE.
only CHOOSE IPV6 bit	single UE IP Address Pool Identity IE	UPF interprets as VLAN to identify the subnet for IPv6 IP address allocation.
only CHOOSE IPV6 bit	more than one UE IP Address Pool Identity IEs	UPF ignores additional occurrences of the UE IP Address Pool Identity IE.
both CHOOSE IPV4 and CHOOSE IPV6 bits	single UE IP Address Pool Identity IE	UPF rejects the Session Establishment Request with a cause code- Mandatory IE incorrect.

When N4 session establishment contains...	And Create PDR IE contains ...	Then ...
both CHOOSE IPV4 and CHOOSE IPV6 bits	two UE IP Address Pool Identity IE	UPF interprets the first and second UE IP Address Pool Identity IEs as the VLANs to identify the subnets for IPv4 and IPv6 IP address allocations respectively.
both CHOOSE IPV4 and CHOOSE IPV6 bits	more than two UE IP Address Pool Identity IE	UPF ignores the additional occurrences of the UE IP Address Pool Identity IE.

Criteria for DHCP configuration validation

Before initiating the DORA or SARR process, UPF performs these validations:

- If either or both CHOOSE IPV4 and CHOOSE IPV6 bits are set, then the DHCP configuration for both IPv4 and IPv6 should not be present. Otherwise, UPF sends the cause code "**Mandatory IE Incorrect**" to SMF.
- If either or both CHOOSE IPV4 and CHOOSE IPV6 bits are set and IPv4 or IPv6 addresses are not present, local APN level configurations **ip address alloc-method dhcp-proxy** and **ipv6 address alloc-method dhcpv6-proxy** should be present for DHCPv4 and DHCPv6, respectively. Otherwise, UPF sends the cause code "**IP Allocation Failure**" to UPF.
- The received VLAN ID should be present for either or both IPv4 and IPv6 in the configuration on UPF.

  If the received VLAN ID is not present in the UPF configuration, then the UPF sends the cause code "**NO RESOURCES AVAILABLE**" to SMF.
- If the DHCP server discovery is disabled and the DHCP server IP address is not configured, but the DHCPv4 service is running, then UPF sends the cause code "**NO RESOURCES AVAILABLE**" to SMF.
- When unicast is enabled and the DHCP server IP addresses are also not configured, but the DHCPv6 service is running, then the UPF sends the cause code "**NO RESOURCES AVAILABLE**" to SMF.
- The Slot/Port should be present and linked to the correct interfaces for IPv4 or IPv6 in the configuration. Otherwise, UPF sends the cause code "**IP Allocation Failure**" to SMF.
- If either or both CHOOSE IPV4 and CHOOSE IPV6 bits are set, then the DHCP configuration for both v4 and v6 must be valid. It is required to start DHCPv4 and DHCPv6 services. Otherwise, the call gets rejected and UPF sends the cause code "**IP Allocation Failure**" to UPF.

Validation of IPv4 and IPv6 addresses

The UPF compares the IPv4 or IPv6 address received from the DHCP server with the chunks allocated by SMF during the DORA or SARR process.

IPv4 validation during DORA process

These stages show the IPv4 validation process during the DORA handshake process:

1. UPF sends the DHCP Discover message to the DHCPv4 server.

2. UPF receives the DHCP Offer message from the DHCPv4 server along with a client IP address, lease time, and server identifier.

   UPF accepts the first Offer message it receives and ignores all the subsequent Offer messages from other DHCPv4 servers.

3. UPF validates the received IP address with the IP pool chunks sent by SMF.

If the IP validation is ...	then...
successful	UPF sends a DHCP Request message to the DHCPv4 server and receives the DHCP Acknowledgement message from the DHCPv4 server.
unsuccessful	UPF discards the Offer message and does not send a DHCP Request message to the DHCPv4 server. The UPF then rejects the session.

4. The UPF receives an Acknowledgment message from the DHCP server. The UPF validates if the IP received in the Offer and the Acknowledge messages are the same. UPF also performs other validations like lease validation, chaddr validation, server identifier validation, etc.

   If the IP values received in the Offer and Acknowledgment messages are the same, UPF considers the IP validation as successful. In this case, the UPF forwards the UE IP to the SMF.

   If there is a mismatch between the IP addresses received in the Offer and Acknowledgment messages, UPF considers the IP validation process as unsuccessful. In this case, the UPF sends a DHCP decline message to the DHCPv4 server and the deletes the session.

   If there is a mismatch in the values of Offer packet and Acknowledgment packet except the IP address validation, then the UPF sends a DHCP release message to the DHCPv4 server.

IPv6 validation during SARR process

For IPv6 sessions, UPF performs IPv6 validation after the SARR process when the **enable ipv6-validation** is configured. These stages show the IPv6 validation process:

1. UPF sends the Solicit message to the DHCPv6 server.
2. UPF receives the Advertise message from the DHCPv6 server.
3. UPF sends the Request message to the DHCPv6 server.
4. UPF receives the Reply message from the DHCPv6 server.
5. UPF validates the received IPv6 address with the IPv6 pool chunks sent by SMF.

If the IP validation is ...	then...
successful	UPF sends the successful establishment response with the received IPv6 address.
unsuccessful	UPF sends DHCP DECLINE message to DHCPv6 server to indicate that the UPF will not use the IPv6 given by DHCPv6 server. Establishment Request is rejected for IPv6 only PDN session. Establishment Request is accepted with cause code partial success for IPv4v6 PDN session.

Enabling DHCP-based IP allocation

These are the necessary steps to enable DHCP-based IP allocation in UPF:

Important While enabling the DHCP feature, it is always recommended to enable the DHCP service on the standby UPF first and then on the active UPF. After making the feature CLI changes, save the changed config at boot config to enable the feature.

If you enable this feature for an existing DNN, the ongoing calls will experience service interruptions. If you enable this feature for a new DNN, there will be no service disruptions, as the new call will always have IP allocation from the DHCP server.

Procedure

Step 1 Start the DHCP service .

Step 2 Configure slot/port to enable discovery of DHCP server.

Start the DHCPv4 or DHCPv6 service

The N4 Establishment contains CHV4 bit or CHV6 message bits and UE IP Address Pool Identity IE to indicate that SMF requires an IP address. To enable the IP allocation through the DHCP server, UPF needs to ensure that the DHCP service is enabled.

Therefore, the UPF uses local APN-level configuration to perform the IP address allocation using the DHCP server.

This task allows you to configure the APN to enable DHCP-proxy based IP allocation.

Procedure

Step 1 Use the command **context** *context_name* to create an instance of the context.

Example:

```
[local]UPF(config)# context ingress
[ingress]UPF(config-ctx)#
```

Step 2 Use the command **apn** *apn_name* to create an instance under the APN configuration mode.

Example:

```
[local]UPF(config)# context ingress
[ingress]UPF(config-ctx)# apn intershat
[ingress]UPF(config-apn)#
```

Note

For the DHCPv4, if the DHCP context name is not configured in the APN, the configured IP context name is used as the DHCP context.

Step 3 Use the commands **ip address alloc-method dhcp-proxy** or **ipv6 address alloc-method dhcpv6-proxy** to define the IP address allocation as DHCP-proxy for IPv4 and IPv6 addresses respectively.

Example:

```
[ingress]UPF(config-apn)# ip address alloc-method dhcp-proxy
[ingress]UPF(config-apn)# ipv6 address alloc-method dhcpv6-proxy
[ingress]UPF(config-apn)#
```

Note
For the DHCPv4, if the DHCP service name is not configured in the APN, the DHCP service from the configured IP context name or the DHCP context is used.

Step 4 Use the command **dhcp service-name** *dhcp_service_name* **| { dhcpv6 service-name** *dhcpv6_service_name* **server-profile** *dhcpv6_server_prof* **client-profile dhcpv6-client-prof }** to specify the DHCP service to be used for IPv4 or IPv6.

Example:

```
[ingress]UPF(config-apn)# dhcp service-name dhcp_service
[ingress]UPF(config-apn)# dhcpv6 service-name dhcpv6-service server-profile
dhcpv6-server-prof client-profile dhcpv6-clientprof
[ingress]UPF(config-apn)# exit
```

Step 5 Use the command **bind address** *dhcp/dhcpv6_ip_address* to bind the UPF interface under the Context Configuration mode. Save and exit the current configuration mode.

Example:

```
[ingress]UPF(config-ctx)# bind address 209.165.201.1
[ingress]UPF(config-ctx)# bind address 2001:DB8::1
[ingress]UPF(config-ctx)# exit
```

This configuration starts the DHCP service and defines the address allocation method as DHCP-proxy.

What to do next

Configure the slot/port to send broadcast messages to the DHCP server. For more details, see the Configure slot/port to enable discovery of DHCP server section.

Configure slot/port to enable discovery of DHCP server

This task helps you configure the slot or port for UPF to send broadcast messages to the DHCP server.

Note If the slot/port is configured, then the UPF sends a broadcast message. However, if the slot/port is not configured, then the DHCP server IP addresses have to be configured. In this case, the UPF sends unicast packets.

Before you begin

Before configuring the slot to enable the discovery of DHCP server, you should enable the DHCP service for IPv4. For more details, see the section.

Procedure

Step 1 Enter the Context configuration mode using the command **context** *context_name* to create an instance.

Example:

```
[local]UPF(config)# context ISP
[egress]UPF(config-ctx)#
```

Step 2 Configure the slot/port through which the UPF connects with the DHCP or DHCPv6 server.

a) **For DHCPv4:** Use the command **dhcp-service** *dhcp_service_name* **dhcp server enable-discover slot/port** *slot_port_number* **dhcp-node-id** *node_id* to configure the Slot/port under the Context configuration mode.

Example:

```
[egress]UPF(config-ctx)# dhcp-service dhcp_service
[egress]UPF(config-dhcp-service)# dhcp server enable-discover slot/port 1/10
dhcp-node-id 5
[egress]UPF(config-dhcp-service)# dhcp client-identifier { msisdn | imsi }
[egress]UPF(config-ctx)#
```

Note

- If SMF sends the same VLAN ID for both IPv4 and IPv6 calls, then the slot and port configuration for IPv4 and IPv6 should be the same for both the calls.
- The VLAN configuration shall be the same on all UPF nodes.

b) **For DHCPv6:** Use the CLI **dhcp-server-profile** *dhcpv6_server_prof* **{ enable | disable } dhcpv6-server-unicast** under the Context configuration mode to enable or disable the server unicast option for DHCPv6 server. Also, use the CLI **dhcp-client-profile** *dhcpv6-client-prof* **{ enable | disable } dhcpv6-client-unicast slot/port***slot_port_number* to enable or disable the client unicast option for DHCPv6 on a specific slot or port.

Example:

```
[egress]UPF(config-ctx)# dhcp-server-profile dhcpv6-server-prof
[egress]UPF(config-dhcp-server-profile)# enable dhcpv6-server-unicast
[egress]UPF(config-dhcp-server-profile)# exit
[egress]UPF(config-ctx)# dhcp-client-profile dhcpv6-client-prof
[egress]UPF(config-dhcp-client-profile)# disable dhcpv6-client-unicast slot/port 1/11
[egress]UPF(config-dhcp-client-profile)# exit
[egress]UPF(config-ctx)#
```

Important

In order to use the CLI **disable dhcpv6-client-unicast**, you must have the "Unicast Address Support for DHCPv6 Messages" license.

Also, to use the CLI **slot/port** *slot_port_number*, you must have UPF license.

For more information on the license, contact your Cisco account representative.

Step 3 Disable or enable IP validation for either or both the received IPv4 and IPv6 addresses.

a) **For IPv4:** Use the command **[no] dhcp ip skip-validation** to skip the validation of IPv4 address received from DHCP server. This is an optional command.

Example:

```
[egress]UPF(config-dhcp-service)# dhcp ip skip-validation
[egress]UPF(config-dhcp-service)#
```

b) **For IPv6:** Use the command **disable ipv6-validation** to disable the IP validation process for IPv6 address received from the DHCP server.

Example:

```
[egress]UPF(config-dhcp-client-profile)# disable ipv6-validation
[egress]UPF(config-dhcp-client-profile)# exit
[egress]UPF(config-ctx)#
```

Step 4 Use the command **dhcp lease-expiration-policy { auto-renew | disconnect }** to configure the action required when the IP address lease expires.

Example:

```
[egress]UPF(config-dhcp-service)# dhcp lease-expiration-policy auto-renew
```

Step 5 Verify the status of DHCP service for IPv4 or IPv6.

a) **For IPv4:** Use the CLI **show dhcp-service all** to verify all the IPv4 DHCP service status.

Example:

```
[local]UPF# show dhcp-service all

Service name:                 dhcp_service
  Context:                      ISP
  Bind:                         Done
  Local IP Address:                100.100.103.20
  Next Hop Address:                None
  DHCP Subnet mask used:           host mask
      MPLS-label:                  None

  Service Status:                  Started
  Retransmission Timeout:          3000 (milli-secs)
  Max Retransmissions:             5
  Lease Time:                      600 (secs)
  Minimum Lease Duration:          600 (secs)
  Maximum Lease Duration:          86400 (secs)
  DHCP Dead Time:                  120 (secs)
  DHCP Dead consecutive Failure:   5
  DHCP T1 Threshold Timer:         50
  DHCP T2 Threshold Timer:         88
  DHCP Client Identifier:          Use MSISDN
  DHCP Server Discovery:        Enabled
      Slot/Port:  1/10
      Node Id:    5
  DHCP Algorithm:                  First Server
  DHCP Servers configured:         None
  DHCP server port: 67
  DHCP server rapid-commit:        disabled
  DHCP client rapid-commit:        disabled
  DHCP server check msg size:      disabled
  DHCP relay agent option:         disabled
  DHCP chaddr validation:          enabled
  DHCP IP address validation:   enabled
```

b) Use the command **show dhcpv6-client-profile { all | name }** *client-profile-name* to verify the status of IPv6 services.

Example:

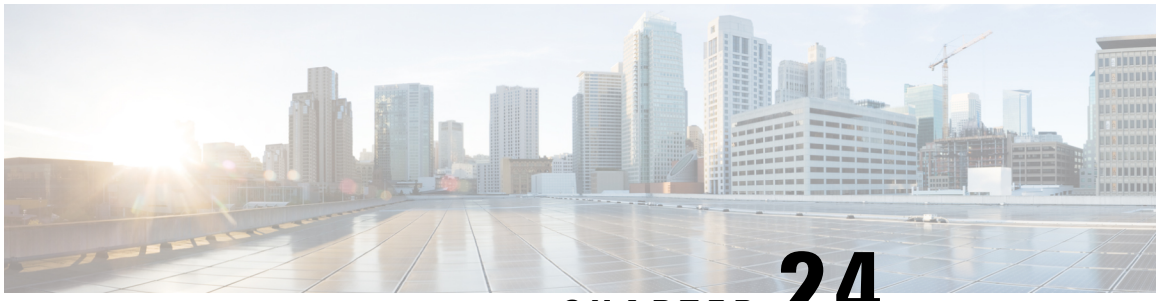
```
[local]UPF# show dhcpv6-client-profile name dhcpv6-client-prof
Client Profile name: dhcpv6-client-prof
Context: ISP
Rapid-commit-dhcpv4: disabled
Rapid-commit-dhcpv6: disabled
Dhcp_msg_spray: disabled
Dns_address: enabled
Netbios_address: disabled
Sip_server_address: enabled
Client_Unicast_to_Server: disabled
IPV6 Validation: disabled
Slot/Port: 1/11
Client identifier: MSISDN
User Class Option: NONE
```

This task enables the discovery of DHCPv4 and DHCPv6 servers through a slot or port number.

Limitations

If an APN is linked to a DHCP service, where the DHCP service is defined in a context different than the context linked to APN, then DORA process does not take place.

For deployments with multiple APNs across different contexts, a separate DHCP service must be configured within each context.



CHAPTER 24

IPsec Support for IPv6

- Feature Summary and Revision History, on page 269
- IPsec AH and ESP, on page 270
- IPsec Transport and Tunnel Mode, on page 270
- IPsec Terminology, on page 270
- Monitoring and Troubleshooting, on page 273

Feature Summary and Revision History

Summary Data

Table 59: Summary Data

Applicable Product (s) or Functional Area	5G-UPF
Applicable Platforms	VPC-SI SMI
Feature Default Setting	Enabled – Always-on
Related Changes in this Release	Not Applicable
Related Documentation	*UCC 5G UPF Configuration and Administration Guide*

Revision History

Table 60: Revision History

Revision Details	Release
First introduced	2021.04.0

Feature Description

IPsec is a suite of protocols that interact with one another to provide secure private communications across IP networks. These protocols allow the system to establish and maintain secure tunnels with peer security gateways. IPsec provides confidentiality, data integrity, access control, and data source authentication to IP datagrams.

IPsec AH and ESP

Authentication Header (AH) and Encapsulating Security Payload (ESP) are the two main wire-level protocols that are used by IPsec. They authenticate (AH) and encrypt-plus-authenticate (ESP) the data flowing over that connection.

- AH is used to authenticate – but not encrypt – IP traffic. Authentication is performed by computing cryptographic hash-based message authentication code over nearly all the fields of the IP packet (excluding those which may be modified in transit, such as TTL or the header checksum), and stores this in a newly added AH header that is sent to the other end. This AH header is injected between the original IP header and the payload.
- ESP provides encryption and optional authentication. It includes header and trailer fields to support the encryption and optional authentication. Encryption for the IP payload is supported in transport mode and for the entire packet in the tunnel mode. Authentication applies to the ESP header and the encrypted data.

IPsec Transport and Tunnel Mode

Transport Mode provides a secure connection between two endpoints as it encapsulates the IP payload. The Tunnel Mode encapsulates the entire IP packet to provide a virtual secure hop between two gateways.

Tunnel Mode forms the more familiar VPN functionality, where entire IP packets are encapsulated inside another and delivered to the destination. It encapsulates the full IP header and the payload.

Note The UPF:UPF ICSR over IPsec works only with Tunnel Mode. Transport Mode is not supported.

IPsec Terminology

Crypto Access Control List

Access Control Lists define rules, usually permissions, for handling subscriber data packets that meet certain criteria. Crypto ACLs, however, define the criteria that must be met for a subscriber data packet to be routed over an IPsec tunnel.

Unlike other ACLs that are applied to interfaces, contexts, or one or more subscribers, crypto ACLs are matched with crypto maps. In addition, crypto ACLs contain only a single rule while other ACL types can consist of multiple rules.

Before routing, the system examines the properties of each subscriber data packet. If the packet properties match the criteria that are specified in the crypto ACL, the system initiates the IPsec policy that is dictated by the crypto map.

Transform Set

Transform Sets are used to define IPsec security associations (SAs). IPsec SAs specify the IPsec protocols to use to protect packets.

Transform sets are used during Phase 2 of IPsec establishment. In this phase, the system and a peer security gateway negotiate one or more transform sets (IPsec SAs) containing the rules for protecting packets. This negotiation ensures that both peers can properly protect and process the packets.

ISAKMP Policy

Internet Security Association Key Management Protocol (ISAKMP) policies are used to define Internet Key Exchange (IKE) SAs. The IKE SAs dictate the shared security parameters (such as which encryption parameters to use, how to authenticate the remote peer, and so on) between the system and a peer security gateway.

During Phase 1 of IPsec establishment, the system and a peer security gateway negotiate IKE SAs. These SAs are used to protect subsequent communications between the peers including the IPsec SA negotiation process.

Crypto Map

Crypto Maps define the tunnel policies that determine how IPsec is implemented for subscriber data packets.

There are several types of crypto maps that are supported in 5G-UPF. They are:

- Manual crypto maps
- IKEv2 crypto maps
- Dynamic crypto maps

Crypto Template

A Crypto Template configures an IKEv2 IPsec policy. It includes most of the IPsec parameters and IKEv2 dynamic parameters for cryptographic, and authentication algorithms. Security gateway service cannot function without a configured crypto template.

Only one crypto template can be configured per service.

Supported Algorithms

IPsec in 5G-UPF supports the protocols in the following table, which are specified in RFC 5996.

Protocol	Type	Supported Options (with VPP)
Internet Key	IKEv2 Encryption	

Protocol	Type	Supported Options (with VPP)
Exchange version 2	IKEv2 Pseudo Random Function	PRF-HMAC-SHA1, PRF-HMAC-MD5, AES-XCBC-PRF-128
	IKEv2 Integrity	HMAC-SHA1-96, HMAC-SHA2-256-128, HMAC-SHA2-384-192. HMAC-SHA2-512-256, HMAC-MD5-96, AES-XCBC-96
	IKEv2 Diffie-Hellman Group	Group 1 (768 bit), Group 2 (1024 bit), Group 5 (1536 bit), Group 14 (2048 bit)
IP Security	IPsec Encapsulating Security Payload Encryption	NULL, DES-CBC, 3DES-CBC, AES-CBC-192, AES-CBC-128, AES-CBC-256, AES-128-GCM-128, AES-128-GCM-64, AES-128-GCM-96, AES-192-GCM, AES-256-GCM-128, AES-256-GCM-64, AES-256-GCM-96
	Extended Sequence Number	Value of 0 or off is supported (ESN itself is not supported)
	IPsec Integrity	NULL, HMAC-SHA1-96, HMAC-MD5-96, HMAC-SHA2-256-128, HMAC-SHA2-384-192, HMAC-SHA2-512-256 **Important** HMAC-SHA2-384-192 and HMAC-SHA2-512-256 are not supported on VPC-DI and VPC-SI platforms if the hardware does not have a crypto hardware.

Limitations and Restrictions

Following are the limitations and restrictions for this feature:

- The feature does not support modification of application ToS.
- If the reordering of packets occurs in an SA, the receiver may discard packets because of anti replay mechanism.
- IPv4 traffic cannot pass through the IPv6 tunnels as this configuration is not allowed. However, IPv4 and IPv6 traffic need IPv4 and IPv6 tunnels respectively.

Example Configurations

Sample Configuration

```
context ipsec-s
ipv6 access-list foo6
permit ip host 2002::1 host 2001::1
#exit
ipsec transform-set B-foo6
#exit
ikev2-ikesa transform-set ikesa-foo6
```

```
#exit
crypto map foo6 ikev2-ipv6
match address foo6
authentication local pre-shared-key encrypted key <encrypted_key>
authentication remote pre-shared-key encrypted key <encrypted_key>
ikev2-ikesa max-retransmission 3
ikev2-ikesa retransmission-timeout 15000
ikev2-ikesa transform-set list ikesa-foo6
ikev2-ikesa rekey
payload foo6-sa0 match ipv6
ipsec transform-set list B-foo6
rekey keepalive
#exit
peer fd4d:5643:2886:6e::7c:1
ikev2-ikesa policy error-notification
#exit
interface ike
ipv6 address fd4d:5643:2886:6e:6b::1/64
crypto-map foo6
#exit
interface loop1 loopback
ipv6 address 2002::1/128
#exit
subscriber default
exit
aaa group default
#exit
ipv6 route 2001::1/128 next-hop fd4d:5643:2886:6e::7c:1 interface ike
#exit
#exit
end
```

Monitoring and Troubleshooting

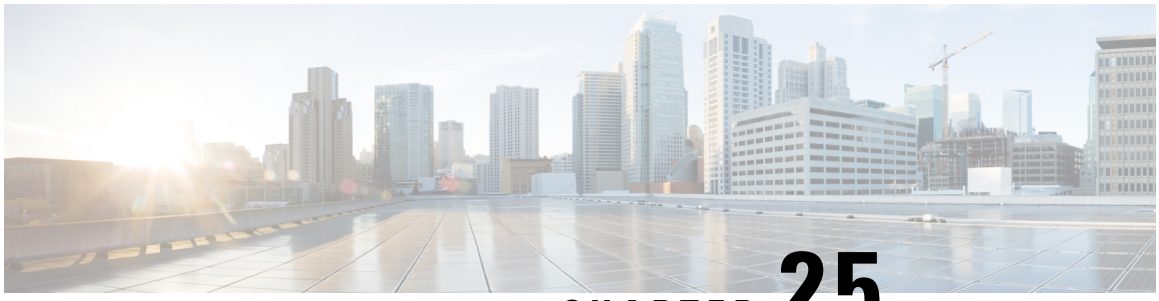
This section describes the CLI commands available to monitor and troubleshoot the IPsec support for the IPv6 feature.

Show Commands

This section provides information about show commands and their outputs in support of this feature.

- **show crypto map**
- **show crypto map tag** *map_name*: Use this command to verify the map status.
- **show crypto map summary**
- **show crypto ikev2-ikesa security-associations**
- **show crypto ikev2-ikesa security-associations tag** *map_name*
- **show crypto ikev2-ikesa security association summary**: Use this command to verify if the IKEv2 SAs are initiated.
- **show crypto ipsec security associations**: Use this command to verify if the IPsec SAs are stabilized.
- **show crypto ipsec security-associations tag** *map_name*
- **show crypto ipsec security-associations peer** *Peer IP Address*

- **show crypto ipsec security-associations summary**
- **show crypto statistics**
- **clear crypto ike-all** — Clear IKEv1 SA / IKEv2 SA of a map based on given criteria.
- **clear crypto managers** — Clear crypto managers.
- **clear crypto statistics** — Clear crypto statistics for this context.
- **clear crypto ikev2 { local-gateway | peer | tag }**
- **clear crypto security-associations { all | counters | local-gateway | peer | tag }**



CHAPTER 25

IPv6 Prefix Delegation

- IPv6 Prefix Delegation, on page 276
- Licensing, on page 276
- Guidelines and Limitations of IPv6 Prefix Delegation, on page 277
- Delegation of IPv6 Prefix, on page 277
- How IPv6 Prefix Delegation Works, on page 279
- Configure DHCP Service Name for Prefix Delegation , on page 281
- Verify IPv6 Prefix Delegation, on page 283

IPv6 Prefix Delegation

Table 61: Feature History

Feature Name	Release Information	Description
IPv6 Prefix Delegation Support to the requesting UE or CPE	2024.04.0	UPF supports IPv6 Prefix Delegation to the requesting User Equipment (UE) or Customer Premises Equipment (CPE). This feature ensures efficient and dynamic allocation of IPv6 prefixes, facilitating seamless connectivity and address management for devices within the network. This feature requires a valid license. Contact your Cisco account representative for more information. **Command Introduced**: The existing CLI **ipv6 address alloc-method no-dynamic allow-prefix-delegation pd-alloc-method no-dynamic** in the APN Configuration mode must be configured to enable this feature. Additionally, DHCP service should also be associated to the APN configuration. **Default Setting**: Disabled – Configuration Required to Enable

The IPv6 Prefix Delegation feature allows a User Equipment (UE) or Customer Premises Equipment (CPE) to obtain an IPv6 prefix, which it can then use to assign IPv6 addresses to devices within its network.

The UPF relies on the SMF for Prefix allocation since the UPF does not manage an IP pool locally. In UPF, the IPv6 Prefix delegation feature is controlled at the DNN (APN) level to ensure that the correct network context is applied.

Licensing

The DHCPv6 PD Prefix & IPv6 Interface-ID via Radius 1K License is required for the IPv6 Prefix Delegation feature. Contact your Cisco account representative for more information.

Guidelines and Limitations of IPv6 Prefix Delegation

The following are the prerequisites for the IPv6 Prefix delegation feature:

- A Packet Data Unit (PDU) session for the UE or CPE must be successfully established, and an IPv6 address must be allocated before a request for IPv6 Prefix Delegation is initiated by UE.
- Support for IPv6 Prefix Delegation is present on SMF.

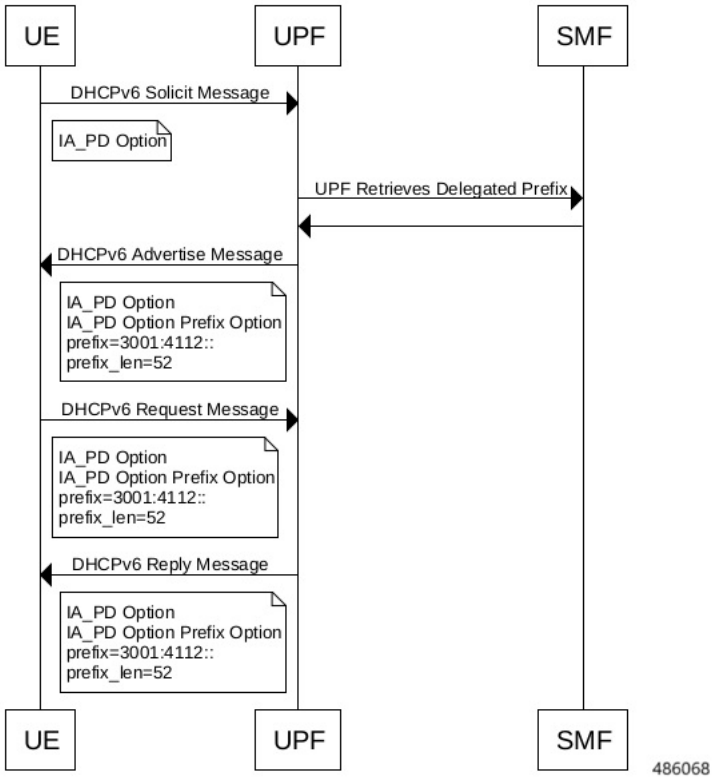
The limitations are:

- Delegated prefix is assigned to UE only if Prefix length received in DHCPv6 Solicit message from UE matches with Prefix length configured on SMF. UE request for a Preferred Prefix sent in Solicit message will not be considered.
- Prefix Exclude Option for DHCPv6-based Prefix (RFC 6603) is not supported.
- Rapid commit option is not supported.
- IPv6 prefix delegation feature will not interwork with IP source violation feature on SxA. Uplink packets coming from devices behind UE will get dropped at SxA.
- UPF doesn't support parsing of retransmitted UE initiated DHCPv6 messages such as Solicit, Request or Release.
- UPF doesn't support DHCPv6 Renew and Rebind procedures as the respective timers are set to infinite.
- UPF doesn't support in-service Downgrade with the feature enabled. MOP should be followed to clear the existing IPv6 prefix delegated session from SMF and UPF and related SMF configurations and IP chunks should be removed before any UPF downgrade.

Delegation of IPv6 Prefix

This section describes the call flow and stages in Delegation of IPv6 prefix.

Figure 15: Delegation of IPv6 Prefix



The following stages describe interactions between SMF (Session Management Function) and UPF (UserPlane Function) for allocating and managing delegated prefixes.

1. The UE/CPE sends DHCPv6 Solicit messages to the UPF. The UE/CPE can include a prefix length inIA_PD Prefix option.
2. UPF receives the requested messages over N3 or S1U/S5U interface on the GTPU tunnel that was created when the PDN session for the UE was established. Since UPF does not manage IP pool locally, UPF acts as a proxy DHCP server and receives the prefix from SMF. SMF provides delegated prefix as per the DNN configuration. SMF rejects the request, if DHCP solicit is received with a prefix length different from the configured prefix length value. If prefix length is not received, SMF allocates as per the configured prefix length
3. UPF sends response to the DHCPv6 Solicit message with a DHCP advertise message containing the delegated prefix in the IA_PD option.

   SMF allocates delegated prefix from the local IPAM pool.

Note When UPF fails to allocate a prefix requested by the UE/CPE, it responds to the DHCPv6 Solicit with an DHCPv6 Advertise message containing IA_PD option with Status Code: NoPrefixAvail.

IPv6 Prefix Release

UPF sends DHCPv6 Reply in response to the DHCPv6 Release message received from UE/CPE confirming the release of prefix address at UPF.

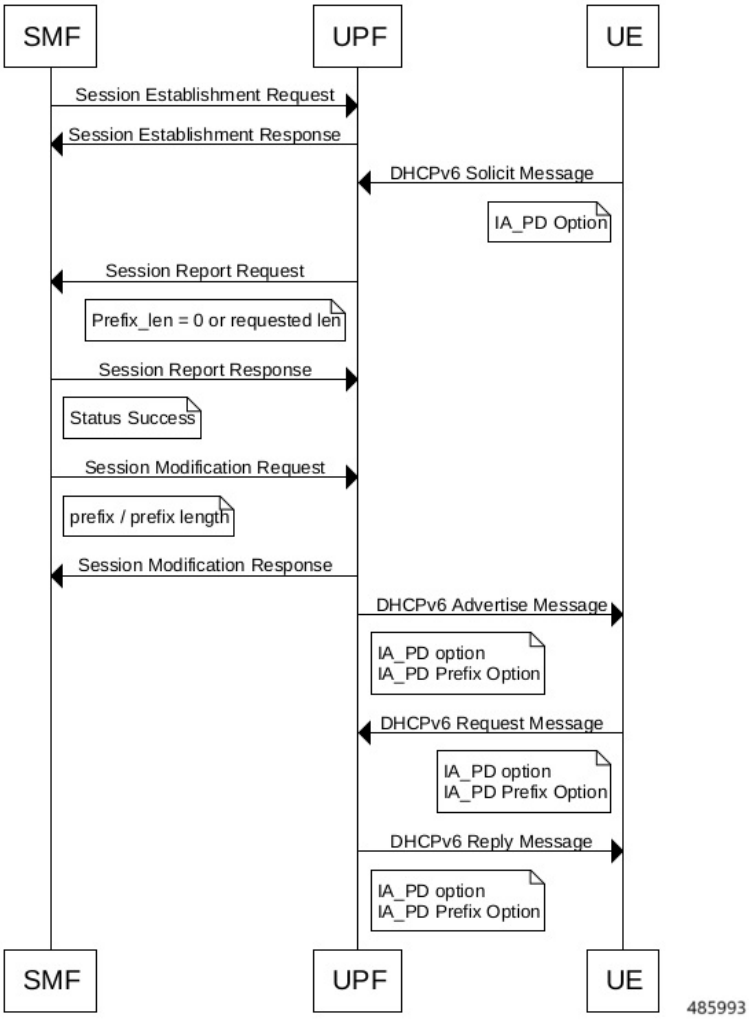
How IPv6 Prefix Delegation Works

The delegated IPv6 prefix is either pushed by the SMF to UPF during session establishment or pulled by UPF post session establishment. UPF allocates the IPv6 Prefix Delegation to the UE / CPE through either of the following flows:

1. Session Establishment and Prefix Request
2. Session Modification Request

Session Modificaiton Request based Prefix Allocaiton

Figure 16: Prefix Received in Session Modification Request Workflow



The following process describes how prefix received from SMF in the Session Modification Request is allocated to UE / CPE.

1. If SMF does not send the delegated prefix during PDN Session Establishment, UPF requests the prefix from SMF by triggering a Session Report Request towards the SMF on receiving a DHCPv6 Solicit message from the UE/CPE. UPF sends the prefix length requested by the UE / CPE in the Session Report or sets it to 0 if no IA_PD Prefix option is received in DHCPv6 Solicit message.
2. If the requested prefix length matches to the prefix length configured on SMF, SMF responds with a successful Session Report Response followed by Session Modification Request with the prefix address and prefix length.
3. UPF sends a DHCPv6 Advertise message containing the delegated prefix address and prefix length in the IA_PD Prefix Options as part of IA_PD Option in response to the DHCPv6 Solicit message to indicate

prefix allocation after receiving Session Modificaiton Request from SMF with prefix address and prefix length.

4. UE/CPE sends DHCPv6 Request message with the received prefix address and prefix length to confirm the allocated prefixes and UPF replies with DHCPv6 Reply message containing same prefix address and prefix length.

Note

- If the UE/CPE requested prefix length does not match with the configured prefix length on SMF, SMF responds with a Session Report Response containing failure cause code.
- If the feature is not enabled on SMF and when UPF sends Session Report Request for prefix delegation, SMF responds with a Session Report Response containing failure cause code.
- When UPF fails to allocate a prefix requested by the UE/CPE, it responds to the DHCPv6 Solicit with an DHCPv6 Advertise message containing IA_PD option with Status Code: NoPrefixAvail.

If..	Then..
If SMF does not send the delegated prefix during PDN Session Establishment	UPF requests the prefix from SMF by triggering a Session Report towards SMF on receiving a DHCP solicit message from the UE/CPE. In the Session report request specify the requested prefix_len. UPF uses DIPR bit and PD-Parameters IE to communicate the requested information in the Session Report Request.
If the requested prefix length matches to the prefix length configured on SMF	SMF responds with a session report response followed by Session Modify Request with the PD_Parameters in the PDI of update PDR for the rulebase PDR.
If prefix allocation is successful on SMF	The PD_Parameters contains a valid Prefix and prefix_length.
If the requested prefix length does not match the configured prefix length on SMF	• SMF responds with a session report response containing failure cause code. • UPF responds to the DHCP Solicit with a DHCP Advertise containing an IA_PD with status code in the IA_PD option as NoPrefixAvail.

Configure DHCP Service Name for Prefix Delegation

Use this procedure to enable DHCP service in server-mode for Prefix delegation.

Procedure

Step 1 Log in to the Context Configuration mode.

Step 2 Enter the **dhcpv6 service-name** *dhcp_service_name* command to enable the DHCP service for Prefix delegation and **bind address** with *valid interface IP | loopback address*. The bind-address for the DHCP service is the IP assigned to a loopback interface configured in the context.

Example:

```
config
context ingress
        dhcpv6-service dhcp_service_name
             bind address {valid interface IP | loopback address }
      exit
```

Step 3 Enter the **apn** *apn_name* command to specify a APN name associated by the UPF node.

Example:

```
config
context ingress
dhcpv6-service dhcp_service_name
             bind address {valid interface IP | loopback address }
      exit
      apn intershat
          ipv6 address alloc-method no-dynamic allow-prefix-delegation pd-alloc-method no-dynamic
      dhcpv6-service dhcp_service_name
exit

exit
```

Step 4 Enter the **ipv6 address alloc-method no-dynamic allow-prefix-delegation pd-alloc-method no-dynamic** command to configure the APN to allow DHCPv6 prefix-delegation.

Example:

```
config
context ingress
dhcpv6-service dhcp_service_name
             bind address {valid interface IP | loopback address }
      exit
      apn intershat
          ipv6 address alloc-method no-dynamic allow-prefix-delegation pd-alloc-method no-dynamic
        dhcpv6-service dhcp_service_name
      exit

exit
```

Step 5 Enter the **exit** command to end the current configuration mode.

Verify IPv6 Prefix Delegation

This section provides information about show commands and their outputs for the IPv6 Prefix delegation feature.

show subscribers user-plane-only full all callid

The **show subscribers user-plane-only full all callid** *call_id* output includes the prefix delegated to the subscriber.

```
…
NAT Policy NAT64: n/a
Converged Session: No    Converged Peer Callid:      n/a
Visited Call: No
Prefix Delegation:
  IPv6 Delegated Prefix: 3001::/48 Sent to UE: Yes
  Preferred Lifetime: 0xFFFFFFFF
  Valid Lifetime: 0xFFFFFFFF
Subscriber Parameters:
IMSI: 404019011037182
IMEI: 1122334455667788
 …
```

Note If the UE does not request the IPv6 Prefix, the show command displays the **Sent to UE: No** value.

If the IPv6 Prefix delegation is not configured or supported on UPF and no Prefix is received from SMF, following output details is displayed.

```
NAT Policy NAT64: n/a
Converged Session: No    Converged Peer Callid:      n/a
Visited Call: No
Prefix Delegation:
  IPv6 Delegated Prefix Not Received
Subscriber Parameters:
IMSI: 404019011037182
IMEI: 1122334455667788 …
```

Note If the UE releases the prefix via DHCPv6 Release, the show command displays the above output.

show dhcpv6 statistics

When APN is configured to receive Delegated Prefix from the local pool, the sessions statistics is visible under the CLI command output of **show dhcpv6 statistics** and displays the following output.

```
DHCPv6 Session Stats:

  Total Current:                  0
     DHCP Proxy:                  0
     DHCP Server:                 0
```

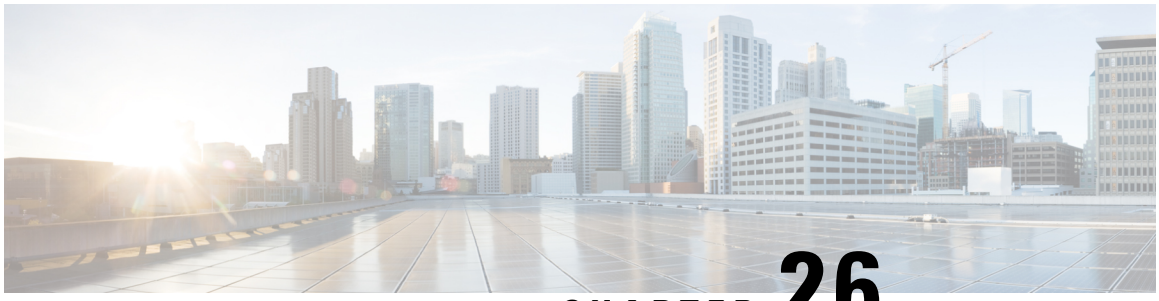
```
   DHCP PD:                        0
   Radius/SMF PD:                  0
   Local PD :                      0

Total Setup:                       0
   DHCP Proxy:                     0
   DHCP Server:                    0
   DHCP PD:                        0
   Radius/SMF PD:                  0
   Local PD :                      0

Total Released:                    0
   DHCP Proxy:                     0
   DHCP Server:                    0
   DHCP PD:                        0
   Radius/SMF PD:                  0
   Local PD :                      0

Session Release Reasons:(dhcp-proxy)
   Admin Releases:                 0   Bearer Call Terminated:         0
   Lease Exp Policy:               0   No Address Available:           0
   IP Address mismatch:            0   Unspecified Failure:            0
   Not on link status:             0   Other Reasons:                  0

Session Release Reasons:(dhcp-prefix-delegation)
   PDNs Released:                  0   Lease Exp Policy:               0
   UE Initiated Release:           0   Other Reasons:                  0
```



CHAPTER 26

LTE - Wi-Fi Seamless Handover

- Feature Summary and Revision History, on page 285
- Feature Description, on page 285
- How It Works, on page 286

Feature Summary and Revision History

Summary Data

Applicable Product(s) or Functional Area	5G-UPF
Applicable Platform(s)	VPC-SI SMI
Feature Default Setting	Enabled - Always-on
Related Changes in this Release	Not Applicable
Related Documentation	Not Applicable

Revision History

Revision Details	Release
First introduced.	2021.02.0

Feature Description

Seamless handovers between LTE and Wi-Fi (S2b), for UEs that need continuity with their ongoing data session, is supported in the 5G UPF architecture.

When handover is initiated from LTE to Wi-Fi, the Delete Bearer Request (DBR) is sent over the LTE tunnel immediately when the Create Session Response (CSR) is sent on the Wi-Fi tunnel. This causes some packet loss because of the IPsec tunnel establishment delay at the ePDG. To address the issue of packet loss, a Delete

Bearer Request is sent on LTE tunnel only on expiry of the configured handover timer. If the LTE tunnel is active, uplink and downlink data is exchanged on the LTE tunnel. When handover is complete, uplink and downlink data is exchanged on the Wi-Fi tunnel. This prevents packet loss. During Wi-Fi to LTE handover, if the Modify Bearer Request is received with HI=1, it initiates a tunnel switch from Wi-Fi to LTE as per the specification.

With this feature, the following benefits are seen:

- Minimum packet loss during LTE to Wi-Fi (S2bGTP) handover and making the handover seamless (that is, MAKE before BREAK).
- LTE procedures are handled gracefully over the LTE tunnel when both tunnels are established with the P-GW.
- Wi-Fi procedures are handled gracefully over the Wi-Fi tunnel when both tunnels are established with the P-GW.



Important

- In an LTE to Wi-Fi or Wi-Fi to LTE handover, a tunnel identifier is allocated for new access traffic type for experiencing seamless handover.

How It Works

EPC to Non-3GPP Untrusted Wi-Fi Handover Call Flow

This section describes the EPC to non-3GPP untrusted Wi-Fi handover call flow.

Figure 17: EPC to Non-3GPP Untrusted Wi-Fi Handover Call Flow

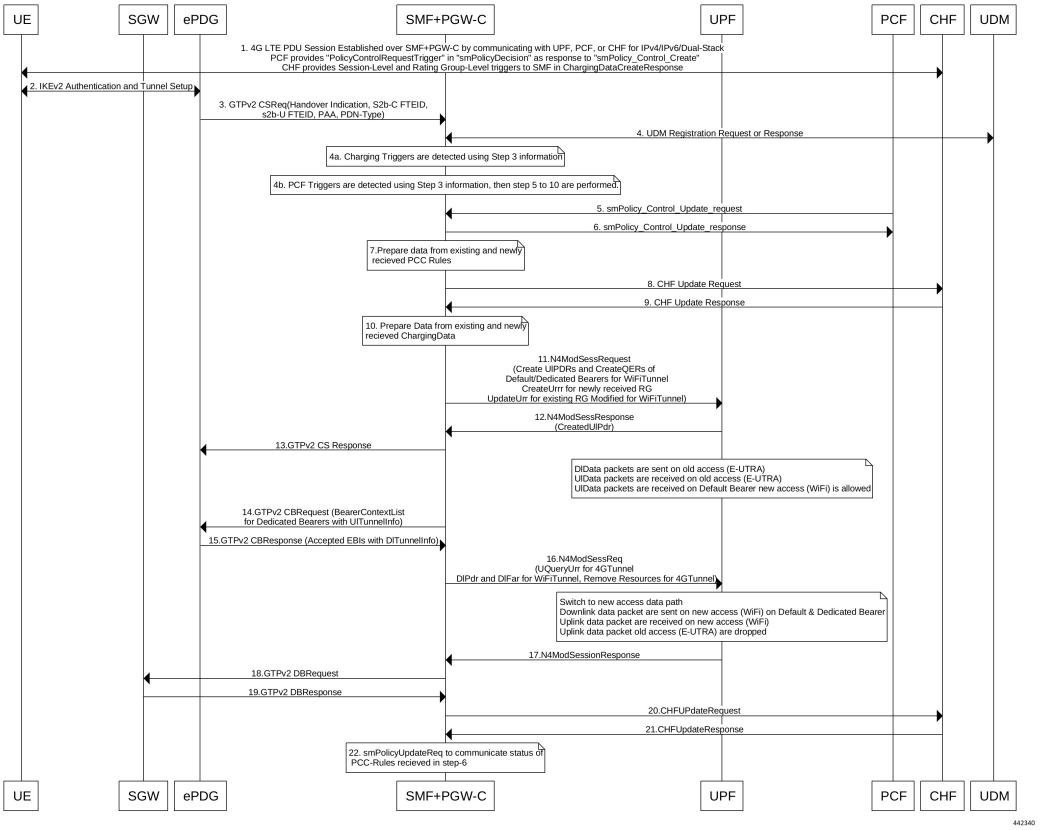


Table 62: EPC to Non-3GPP Untrusted Wi-Fi Handover Call Flow Description

Step	Description
1	The UE is attached to the 3GPP access network. The SMF+PGW-C communicates with UPF, PCF, and CHF for IPv4, IPv6, or dual-stack to establish 4G LTE PDU session. The PCF sends the Policy Control Request trigger, which is the SM policy decision, in response to SM policy control create. The CHF provides session-level or rating-group-level triggers to the SMF in Charging Data Create response.
2	The UE connects to an untrusted non-3GPP access and an ePDG is selected through the ePDG selection process. Then, the UE initiates the handover attach procedure as defined in *3GPP TS 23.402, section 8.6.2.1*. After the IKE tunnel is established between the UE and ePDG and after the UE is authenticated over SWm interface with AAA server, the UE initiates IKE authentication (IKE_AUTH). The IKE_AUTH includes configuration parameters of the earlier assigned IPv4 or IPv6 addresses in the EPC and P-CSCF and the DNS options.

Step	Description
3	The ePDG sends a Create Session Request to the P-GW. This request includes the following details: • IMSI • APN • Handover indication • RAT type • ePDG TEID of the Control Plane • ePDG address for the User Plane • ePDG TEID of the User Plane • EPS bearer identity • User location The RAT type indicates the non-3GPP access technology type. If the UE supports the IP address preservation and is included in the port analyzer adapter (PAA), then the ePDG configures the handover indication in the Create Session Request to allow the PDN gateway to reallocate the same IP address or the prefix assigned to the UE. This IP address or prefix is assigned while UE is connected to the 3GPP IP access and initiates the policy modification procedure with PCF.
4a	The SMF performs UDM registration by updating the PGW-C FQDN with UDM. The UDM registration does not occur during the session establishment with EPC.
4b	The SMF detects the charging triggers with the information available in Step 3 against the charging triggers that are received during EPC session establishment.
4c	The SMF detects the PCF triggers with the information available in Step 3 against the Request Policy Control triggers that are received in the communication with PCF during EPC session establishment.
5	Based on the detected armed Policy Control Triggers that are received in Step 4b, the SMF sends the SM Policy Control Update request with the detected access parameters in Step 3 to the PCF.
6	The PCF includes new or updated PCC rules and sends the SM Policy Control Update response. The Update response includes information on the SM policy decision.
7	Based on the information received in Step 6 and existing policy data of EPC session, SMF prepares the information for the new or updated PCC rules.
8	If new PCC rules are received in Step 6 with new Rating Group that requires quota information, SMF sends the Charging Update request to CHF. SMF also includes new access parameters for the PDU session information.
9	CHF sends the Charging Update Response with multi-unit information that contains quota information for the requested rating-group in Step 8 to SMF. CHF may also send the new quota information for the existing rating-group of EPC session.
10	SMF processes the information that is received as Charging Update response from CHF.

Step	Description
11	SMF sends the N4 session modification request to UPF for Wi-Fi tunnel. This request includes details on creation of uplink PDR, creation of QER, creation of URR for received new rating-group quota information, and update on URR for modified quota information.
12	UPF sends the UL tunnel information that is in created PDR as the N4 session modification response to SMF.
13	SMF sends the GTPv2 Create Session response to S-GW. This response details on request accepted or request accepted partially, P-GW S2b F-TEID, PAA, APN-AMBR, bearer context creation, charging gateway address, and APCO.
14	SMF sends the GTPv2 Create Bearer request to S-GW. This request includes information on bearer context list, which contains DL tunnel information to end-user, to be created.
15	S-GW sends the GTPv2 Create Bearer response to SMF. The response includes details on request accepted or request accepted partially and bearer contexts.
16	SMF processes the Create Bearer response and derives the DL tunnel Information for the established bearer and the the failed EBI list, if any. SMF sends the N4 session modification request to UPF for Wi-Fi tunnel. This request is to create the DL PDR and DL FAR with DL tunnel information for each bearer, RAT modification information, and to delete resources for the 4G tunnel. SMF also deletes the N4 resources of Wi-Fi tunnel for the received failed EBI list or the failed QFI list.
17	UPF sends the usage report as N4 Session Modification response to SMF.
18	SMF+PGW-C sends the GTPv2 DB request to S-GW. This request includes EBI or list of EBIs.
19	S-GW sends the GTPv2 DB response to SMF+PGW-C.
20	SMF sends the Charging Update request to CHF. This request includes the PDU session information with the new access params and multi-usage report containing details on the access params and usage report that is received in Step 8
21	CHF sends the multi-unit information as Charging Update response to SMF. The multi-unit information may include new quota information for the existing rating-groups.
22	SMF sends the SM Policy Control Update request to UPF. This request includes the new access params and rule report for failed QFI list that is received from AMF as part of Create Bearer response. PCF sends the SM policy decision as SM Policy Control Update response. SMF processes the SM policy decision and handles it as PCF Initiation Modify procedure as defined in *3GPP 23.502, section 4.3.3.2*.

Non-3GPP Untrusted Wi-Fi to EPC Handover Call Flow

This section describes the non-3GPP untrusted Wi-Fi to EPC handover call flow.

Figure 18: Non-3GPP Untrusted Wi-Fi to EPC Handover Call Flow

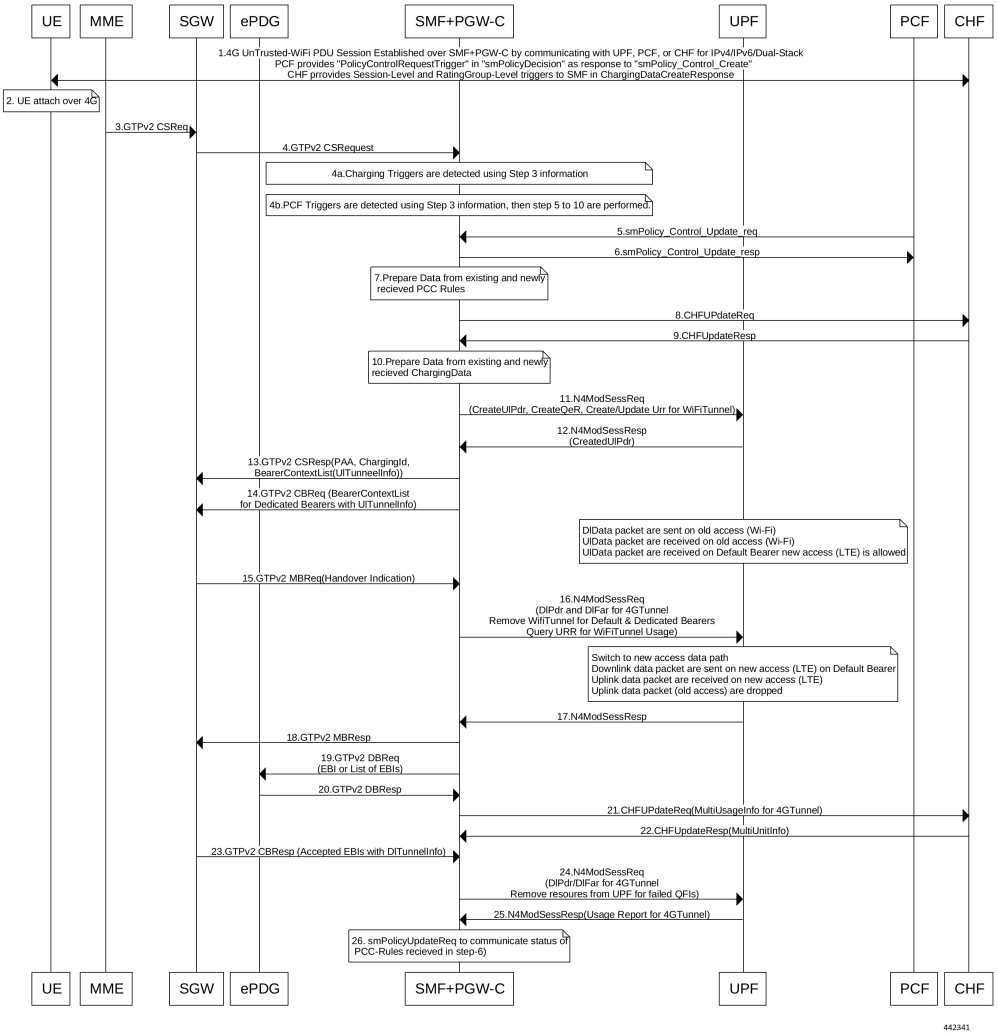


Table 63: Non-3GPP Untrusted Wi-Fi to EPC Handover Call Flow Description

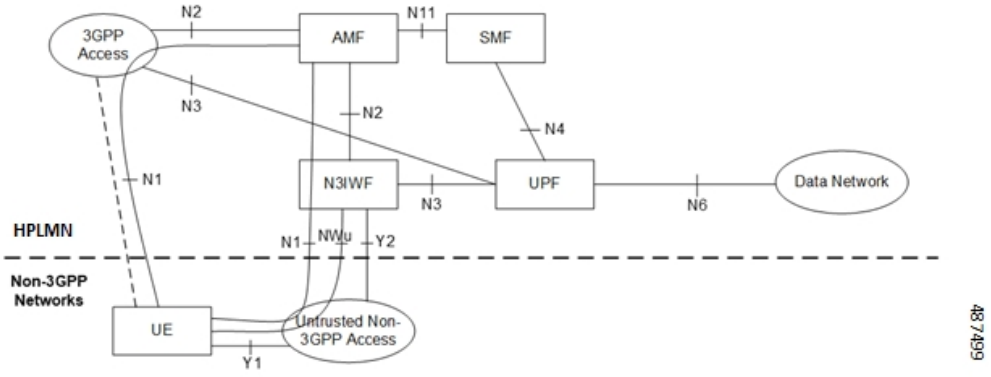
Step	Description
1	One or more PDU sessions are established between UE and ePDG through untrusted non-3GPP access. With the 5G NAS capability of UE, ePDG selects a combined PGW+SMF. UE sends the PDU session ID to the PGW+SMF.
2	UE discovers the E-UTRAN access and hands over the sessions from the currently used non-3GPP access system to E-UTRAN. For details on UE discovery of the 3GPP access system, see *3GPP TS 23.401*, section 4.8. UE sends an Attach request to MME for the Handover Attach request type. E-UTRAN routes the messages received from UE to MME as defined in *3GPP TS 23.401*. UE includes the one of the APNs which are corresponding to the PDN connections in the source non-3GPP access. The APN is provided as defined in *3GPP TS 23.401*.

Step	Description
3	MME and HSS perform authentication, which is followed by location update procedure and subscriber data retrieval to receive the APN information. The MME selects an APN, an SGW and PDN gateway as defined in *3GPP TS 23.401*. MME sends a Create Session Request message to SGW. This request includes information on IMSI, MME context ID, PDN-GW address, handover indication for the “handover” request type, and APN.
4	SGW sends a Create Session Request, which is handover indication, message to PDN-GW in the HPLMN as described in *3GPP TS 23.401*. As the MME includes the handover indication information in the Create Session Request message, the SGW sends the GTPv2 Create Session Request message to PDN GW. This message includes details on IMSI, APN, handover indication, RAT type, S5-C TEID, S5-U TEID of the user plane, EBI, and user location information. The RAT type indicates the 3GPP IP access E-UTRAN technology type. If the UE supports IP address preservation and is included in PAA, the SGW configures the handover indication in the Creation Session Request. With this configuration, the PDN GW re-allocates the same IP address or prefix that was assigned to the UE while it was connected to the 3GPP IP access. With this configuration, SGW initiates the Policy Modification Procedure to the PCF. As the handover indication is includes, the PDN GW does not switch the tunnel from non-3GPP IP access to 3GPP access system at this point. SMF does not perform the UDM Registration as the registration happens during the Wi-Fi session establishment.
4a	SMF detects the charging triggers with the information available in Step 3 against the charging triggers that are received during EPC session establishment.
4b	SMF detects the PCF triggers with the information available in Step 3 against the Request Policy Control triggers that are received in the communication with PCF during EPC session establishment.
5	Based on the detected armed Policy Control Triggers that are received in Step 4b, SMF sends the SM Policy Control Update request with the detected access parameters in Step 3 to PCF.
6	PCF sends the SM Policy Control Update response, which is the SM policy decision, by including new or updated PCC rules.
7	Based on the information received in Step 6 and existing policy data of EPC session, SMF prepares the information for the new or updated PCC rules.
8	If SMF receives new PCC rules in Step 6, the SMF sends the Charging Update request, with the new rating-group having quota information, to CHF. This request includes the PDU session information with the new access params.
9	CHF sends the multi-unit information as Charging Update response to SMF. The multi-unit information includes new quota information for the rating-group and the existing rating-group of EPC session, if any.
10	SMF prepares the charging data of the received Charging Update Response that CHF sent.
11	SMF sends the N4 Session Modification Request to UPF. This request includes the details on creation of UL and DL PDR, creation of QER, creation of URR for received new rating-group quota information, updated URR for modified quota information, and creation of FAR.
12	UPF sends the UL tunnel information in the created PDR as N4 Session Modification response to SMF.

Step	Description
13	SMF sends the GTPv2 Create Session response to S-GW. This response details on request accepted or request accepted partially, P-GW S2b F-TEID, PAA, APN-AMBR, bearer context creation, charging gateway address, and APCO.
14	SGW sends the Modification Bearer request with handover indication to PGW for data path switching from Wi-Fi tunnel to 4G tunnel.
15	PGW sends the N4 Session Modification request to delete the Wi-Fi tunnel and to configure DL tunnel information that is received in GTPv2 Create Session request for 4G tunnel in Step 4.
16	UPF sends the N4 Session Modification response to SMF.
17	SMF sends the GTPv2 Create Session request, which includes the bearer context list, to SGW. This list includes the DL Tunnel information for the end-user.
18	SGW sends the GTPv2 Create Session response to SMF. This response includes details on request accepted or request accepted partially and bearer contexts.
19	ePDG sends the GTPv2 Create Bearer resp (accepted EBIs with DL tunnel info to SMF
20	SMF processes the Create Bearer response and derives the DL tunnel Information for the established bearer and the failed EBI list, if any. SMF sends the N4 session modification request to UPF for Wi-Fi tunnel. This request is to update the DL FAR with the DL tunnel information, RAT modification information, and to delete resources for the 4G tunnel. SMF also deletes the N4 resources of Wi-Fi tunnel for the received failed EBI list or the failed QFI list.
21	UPF sends the N4 Session Modification Response with usage report to SMF.
22	SMF sends the Charging Update request to CHF. This request includes the PDU session information with new access params and multi-usage report consisting of access-params and usage report that is received in Step 8.
23	CHF sends the Charging Update Response with multi-unit information that contains quota information for the existing rating-groups to SMF.
24	SMF+PGW-C initiates the GTPv2 DB Request toward SGW by including EBI or EBI list.
25	SGW sends the GTPv2 DB Response toward SMF+PGW-C.
26	SMF sends the SM Policy Control Update request to UPF. This request includes the new access params and rule report for failed QFI list that is received from AMF as part of Create Bearer response. PCF sends the SM policy decision as SM Policy Control Update response. SMF processes the SM policy decision and handles it as PCF Initiation Modify procedure as defined in 3GPP 23.502 section 4.3.3.2.

Non-3GPP Interworking Function

Starting with 3GPP Release 15, the Non-3GPP Interworking Function (N3IWF) facilitates non-3GPP access to the core network by adapting untrusted Wi-Fi (802.11) for integration with the 5G Core. N3IWF in the UPF manages traffic seamlessly from both trusted 3GPP RAN and untrusted Wi-Fi access points, such as public hotspots or private routers. It ensures secure end-to-end communication, independent of link-layer security like WPA2. The UPF anchors traffic from both access types, enabling reliable and scalable connectivity across networks.



Refer to N3IWF for Non-3GPP Access for N3IWF attach and handover procedure. In UPF, this functionality is same as ePDG, except that UPF receives ULI with value NRTAI when call is from N3IWF component. This NRTAI value is populated in existing PFCP_IE_SUB_PARAMS IE in ULI. UPF dumps this NRTAI value in EDR.

ICSR and Session Recovery

- At Session Management Function, during transition, the most recent is considered as the stable state and a full checkpoint is triggered once handover is complete from LTE to Wi-Fi (S2BGTP) or vice-versa. This is applicable to Session Recovery and ICSR. User Plane has individual session recovery and ICSR check pointing on every message received.
- During handover failure, that is, when SMF and UPF are out of sync, the SMF session is recovered on the most recently accessed state and UPF is recovered in the new transition state. This behavior is applicable during UPF failure.

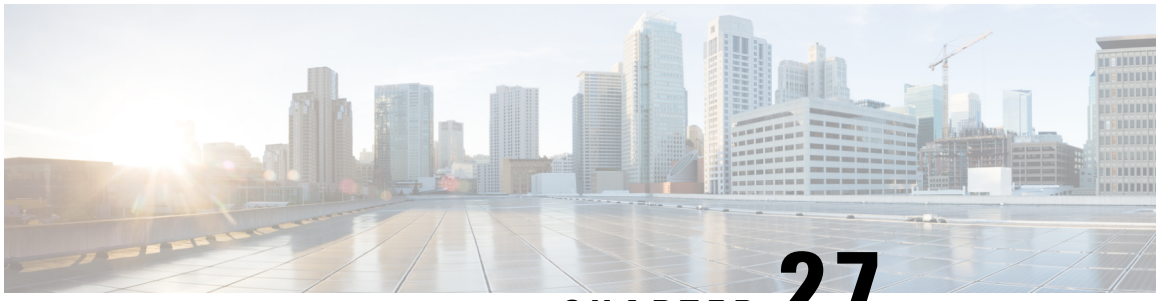
Limitations

The LTE - Wi-Fi Seamless Handover feature does not support LTE to eHRPD and Wi-Fi to eHRPD handover and hand back.

Standards Compliance

The LTE – Wi-Fi Seamless Handover feature is compliant with the following standards:

- 3GPP TS 23.214
- 3GPP TS 29.244
- 3GPP TS 23.401
- 3GPP TS 23.402



CHAPTER 27

Monitor Subscriber

- Feature Summary and Revision History, on page 295
- Feature Description, on page 296
- How It Works, on page 297
- Configuring the Hexdump Module for MonSub in UPF, on page 311
- Monitoring and Troubleshooting, on page 312

Feature Summary and Revision History

Summary Data

Table 64: Summary Data

Applicable Product (s) or Functional Area	5G-UPF
Applicable Platforms	VPC-SI SMI
Feature Default Setting	Disabled – Configuration Required
Related Changes in this Release	Not Applicable
Related Documentation	*UCC 5G UPF Configuration and Administration Guide*

Revision History

Revision Details	Release
Added support for subscriber tracing from SMF.	2023.02.0
In this release, enhancement related to MonSub CLI, subscriber tracing limits, packet processing throughput, PCAP success and, error code notifications are added to this feature. The CP and UP session manager functionality is also added for this release.	21.22.x

Revision Details	Release
First introduced.	21.16.1

Feature Description

The Monitor Subscriber (MonSub) feature enables tracing of subscriber-related information which includes user and control traffic, and events such as charging and internal events that are useful for debugging. By default, this information is displayed on the Session Management Function (SMF) console, where the user executes MonSub tracing CLI command and captured in a Packet Capture (PCAP) file on the UPF.

User traffic is carried on slowpath when packets traverse to the application, or fastpath when packets don't have to traverse up to the application and are offloaded to fastpath processing (VPP). Slowpath mode was the default mode until fastpath offload (VPP) into SAEGW was introduced.

Monitor Subscriber provides the following functionality:

- Monitoring user traffic for one subscriber at a time.
- Continuous capture of user traffic from fastpath in PCAP files in the UPF.
- UPF captures both Control and Data packets.
- The control event traffic and other related information are captured in separate PCAP files in the UPF.
- New option, UP PCAP trace [W - UP PCAP Trace (ON)], is introduced for SMF and UPF in MonSub CLI. The slow-path and fast-path PCAP generates only when this option is ON.
- Maximum of 4 MonSub sessions can be enabled per UPF.
- By default, UPF supports the following MonSub options:
  - PFCP Data (Option 49)
  - CSS Data (Option 34)
  - L3 Data (Option 19)
  - GTPU (Option 26)
  - EDR (Option 77)
  - Subscriber summary after call disconnect
- CSS data now additionally captures Flow Termination due to FLOW IDLE TIMEOUT.

Following are some of the important definitions related to this feature:

- **Chassis Traffic Volume**: The total volume of packet throughput on the chassis.
- **Monitored Traffic Volume**: Monitoring of the total throughput of all the subscribers through MonSub across all the MonSub sessions.
- **PCAP Success**: The percentage of the MonSub traffic capture request and the successful capture in the PCAP files.

Packet Processing Throughput

Following are the scenarios that impact the packet processing throughput:

- When VPP utilization is above 80%, MonSub may have an impact to packet processing throughput. The impact is in proportion to the monitored traffic volume.
- Specifically, when the monitored traffic volume approaches 10% of the chassis traffic volume, there may be impact on the VPP throughput causing subscriber packet loss.
- The impact to packet processing throughput is higher when using monitor priorities above 0 (zero).



Caution You must be cautious during the packet processing. When VPP is running at 80% utilization and handling approximately 10Gbps chassis traffic volume, there's an impact on the packet processing, if the set of MonSub sessions is collectively monitoring the subscribers, totaling more than 1Gbps of monitored traffic volume.

PCAP Success

The PCAP success depends on the following factors:

- The level of PCAP success depends on several factors, including monitored traffic volume, VPP utilization, MonSub monitor priority, and background disk I/O.
- In general, the PCAP success rates are greater for the following cases:
  - When the VPP utilization is low and/or MonSub monitor priority is above best-effort.
  - When the monitored traffic volume is less than 10% of the chassis traffic volume.

Example: When VPP is running at 80% utilization and handling approximately 10Gbps chassis traffic volume, then monitored traffic volume up to 1Gbps is likely to yield high PCAP success percentages.

How It Works

The Monitor Subscriber feature is discussed in detail in the following sections:

SMF and UPF Interactions

UPF supports subscriber monitoring using CLI configurations to enable monitor subscriber per UE. However, when there are multiple UPFs in the network, it requires the users to enable monitor subscriber for a UE in all the UPFs, or wait for the UE to attach and identify the UPF, which is handling it and enable monitor subscriber on that UPF.

This feature enables monitor subscriber on UPF based on the configurations on SMF, thereby, eliminates the risk of missing captures.

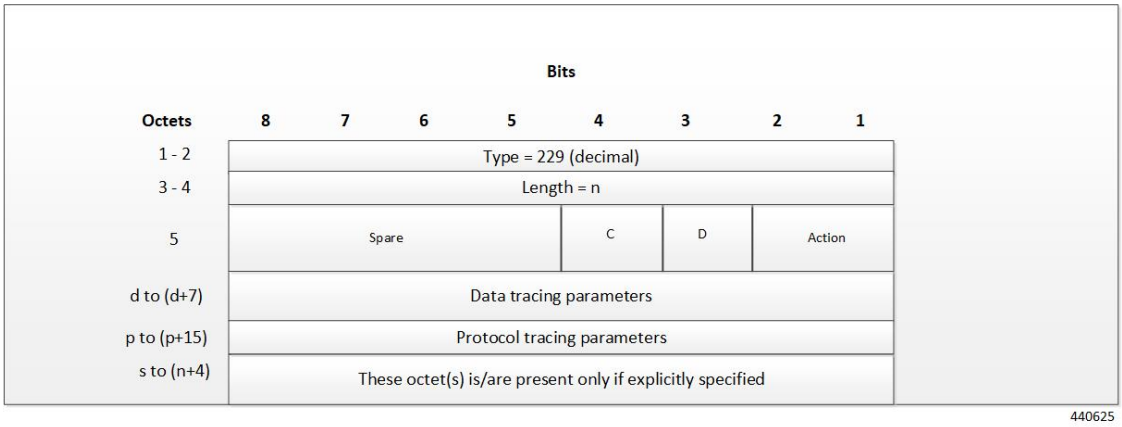
Subscriber Tracing from SMF

MON SUB IE for Subscriber Tracing from SMF

SMF sends a proprietary MON SUB IE to the UPF to start or stop a subscriber tracing session for a given IMSI. The MON SUB IE is a private Information Element (IE), which can be included only when the tracing options selected require U-Plane participation.

PFCP_IE_MONITOR_SUBSCRIBER_INFO

Figure 19: Subscriber Tracing



Action: STOP / START monitor subscriber tracing. STOP =1, START =2.

Note D = DATA events tracing is ON when D=1. The 8 octets (d to d+7) contain data events tracing (fastpath) information should be present only when, D=1.

C = CONTROL events tracing is ON when C=1.

Data Tracing (fastpath) Information (8 octets): It contains the data filter parameters like Packet capture, Packet capture size, and MEH header.

- **Octet 1:**
  - Bit 1: VPP enable/disable
  - Bit 2: FCAP - Packet capture
  - Bit 3: MEH present
  - Bit 4 to 6: Priority
- **Octet 2 to 3:** Packet size
- **Octet 4 – 8:** Reserved for future use. Currently, all set to 0.

Protocol Tracing Information (16 octets/128 bits): The 16 octets, p-to-(p+15), contain protocol tracing information. It is present only when either the control flag (C) or the data flag (D) is enabled. Each bit represents a unique protocol to monitor. For example, if the 49th bit is 1, PFCP events tracing is ON. The control event

flag controls the Protocol Tracing *Rulematch Events (Option 34)*, *L3 Data (Option 19)*, *GTPU (Option 26)*, *EDR (Option 77)*, and *Subscriber Summary After Call Disconnect*.

Enabling Subscriber Tracing from SMF

The SMF sends the MON SUB IE in either of two situations:

1. Session Establishment Request (for a new session), or
2. Session Modification Request (for ongoing session)

Enabling subscriber tracing on SMF is done through a new CLI. For more details, see SMF Monitor Subscriber Configuration.



Note Enabling of session tracing is not piggybacked on any PFCP message. Session Tracing is enabled on UPF via a separate Sx Modification Request from SMF.

Disabling Subscriber Tracing from SMF

To disable the subscriber tracing from SMF, the MON SUB IE with Action = STOP, is sent in a Session Modification Request.

PCAP Generation

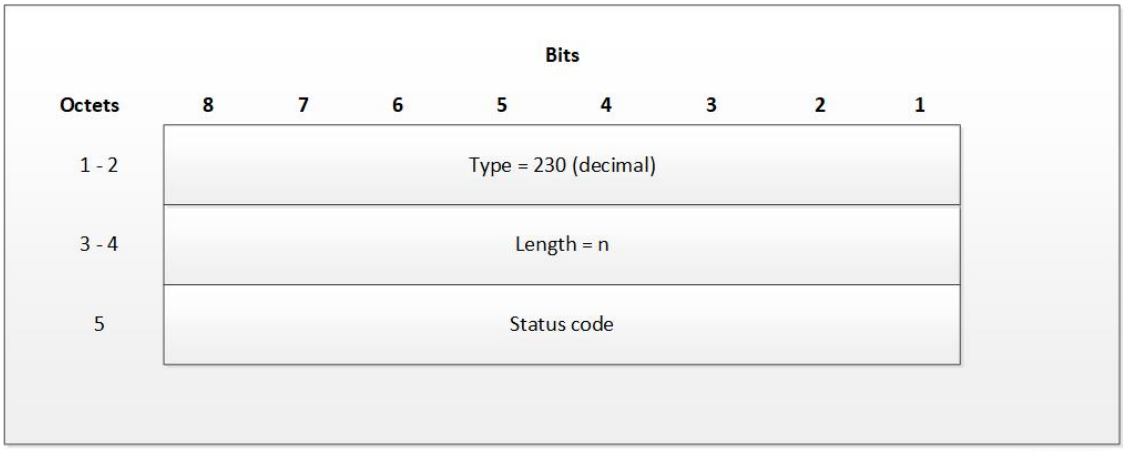
UPF generates the slowpath and fastpath PCAPs in the following 3 situations:

1. When subscriber tracing is stopped from SMF via the N4 interface.
2. When a subscriber session is deleted gracefully.
3. When a subscriber session is locally purged on UPF.

Subscriber Trace Status Report IE (Private IE)

The Subscriber Trace Status Report IE is sent corresponding to all monsub START request received from SMF.

Figure 20: Subscriber Trace Status Report



The status code indicates the acceptance or the rejection of the subscriber trace at UP. Status code = 0 means, a success. Values 1-255 uniquely specifies a specific error code or a notification. The list of error codes and notifications are defined post development.

Table 65: Error Code and Notification Table

Status Code	Status Description
MONSUB_SM_SUCCESS (0)	
MONSUB_SM_ERROR_FAILURE (1)	MonSub: Generic Failure status received
MONSUB_SM_ERROR_UNSUPPORTED (2)	MonSub: Unsupported Failure!
MONSUB_SM_ERROR_SESSION_EXIST_NONE (3)	MonSub: Session not Found!");
MONSUB_SM_ERROR_SESSION_LIMIT_EXCEED (4)	MonSub: Max Connections reached!
MONSUB_SM_ERROR_SESSION_INVALID_PARAM (5)	MonSub: Connect Message Failed!
MONSUB_SM_ERROR_SESSION_ALLOC_FAIL (6)	MonSub: Could not allocate monsub session at NPU!
MONSUB_SM_ERROR_CONFIG_INVALID_PARAM (7)	MonSub: Config Message Failed!
MONSUB_SM_ERROR_MONITOR_LIMIT_EXCEED (8)	MonSub: Max Stream Limit reached!
MONSUB_SM_ERROR_MONITOR_INVALID_PARAM (9)	MonSub: Monitor Message Failed!
MONSUB_SM_ERROR_MAX (10)	MonSub: Max Error!
MONSUB_COPROCDATA_CORRUPTED (11)	MonSub: File Handling Process Failed!
MONSUB_MAX_TRACING_SESSIONS_REACHED (12)	MonSub: Maximum Number of Tracing Sessions reached!

Status Code	Status Description
MONSUB_STOP_RECVD_WAIT_POLL_TIMEOUT (13)	MonSub: STOP notification is Successful. Wait till the poll-timeout configuration to start the next tracing!
MONSUB_FILECOPY_SOURCE_DIR_NOT_EXIST (14)	MonSub: Source Directory does not exist!
MONSUB_FILECOPY_DEST_DIR_NOT_EXIST (15)	MonSub: Destination Directory does not exist!
MONSUB_FILECOPY_SOURCE_DIR_OPEN_FAILURE (16)	MonSub: Unable to open Source Directory!
MONSUB_FILECOPY_DEST_DIR_OPEN_FAILURE (17)	MonSub: Unable to open Destination Directory!
MONSUB_FILECOPY_SOURCE_OPEN_FAILED (18)	MonSub: Unable to open Source File!
MONSUB_FILECOPY_DESTINATION_OPEN_FAILED (19)	MonSub: Unable to open Destination File!
MONSUB_FILECOPY_DONE_FILE_DELETION_FAILED (20)	MonSub: Unable to delete .done file in Source Path!
MONSUB_FILECOPY_PCAP_FILE_DELETION_FAILED (21)	MonSub: Unable to delete .pcap file in Destination Path!
MONSUB_RESPONSE_NPUMGR_MONSUB_SESS_FAILED (23)	MonSub: Messenger Failure during Session Notification to NPUMGR!
MONSUB_RESPONSE_NPUMGR_MONSUB_CFG_FAILED (24)	MonSub-Config push to NPUMGR!
MONSUB_RESPONSE_NPUMGR_MONSUB_MONITOR_FAILED (25)	MonSub-Monitor Notification to NPUMGR!
MONSUB_RESPONSE_COPROC_FAILED (27)	MonSub: File Handling Process Failed!
MONSUB_RESPONSE_FILE_TRANSFER_SUCCESS (28)	MonSub: File Transfer successful.
MONSUB_RESPONSE_FILE_TRANSFER_FAILED (29)	MonSub: File Transfer failed!
MONSUB_ADMINISTRATIVE_DISCONNECT (30)	MonSub: Administrative Disconnect!
MONSUB_FILECOPY_DESTINATION_DISK_FULL (31)	MonSub: No space left in Destination Path!
MONSUB_FILECOPY_COPROC_ABRUPTLY_KILLED (32)	MonSub: File copy co-proc terminated abruptly!
MONSUB_LOGGING_COPROC_ABRUPTLY_KILLED (33)	MonSub: Logging co-proc terminated abruptly!
MONSUB_SM_DISCONNECT (34)	
MONSUB_FILECOPY_STATUS_MAX (35)	

Status Code	Status Description
Other	Internal Error adding protocol monitor trace - aborting…

Fastpath and Slowpath Packet Captures

The packet capture in MonSub takes place in Slowpath and Fastpath.

Slowpath

- Slowpath PCAP captures the packets going through slowpath (application).

  When option PFCP(49), GTPU(26), L3(19), and CSS(34) is enabled, N4 messages/data packets are shown from the time MonSub is enabled.

- Uplink data packet is shown as incoming GTPU, rulematch, outgoing L3.
- Downlink data packet is shown as incoming L3, rulematch, outgoing L3.

Fastpath

- Fastpath PCAP will capture the packets going through fastpath.
- For pure Fastpath packets, the uplink data packet is shown as incoming GTPU, outgoing L3.
- Downlink data packet is shown as incoming L3, outgoing GTPU.
- Slowpath packets are also shown in Fastpath PCAP, while egressing the UPF in both directions.

UPF SessMgr Functionality

Following are the modifications done in the UPF session manager to support this feature:

- Provide services to the CLI for enabling or disabling the MonSub tracing.
- Control NPUMgr to connect/start/stop/add/delete MonSub sessions.
- Based on the instructions from the CLI, configures panopticon (through NPUMgr) for changes, such as packet size and priority.
- Read the "hex dump module" configurations to manage the storage of generated pcap files.
- If the file copy fails or there are problems with session manager MonSub instantiation, raises the SNMP alarms.
- Capture the control or slowpath packets as a separate slowpath PCAP.
- This feature supports a maximum of four monitor subscriber tracing sessions for a UPF instance. The NPUMgr enforces the tracing limit.
- The MonSub tracing session terminates in the absence of no space on hard disk or no hard disk.
- Flow termination due to flow idle timeout is also captured in option 34 CSS Data.
- N4/Sx Establishment Request/Response is also captured.

Multi PDN Multi-trace

For a multi-PDN call, when you start the MonSub with Multi-trace=OFF, then it traces the only one PDN as a part of that MonSub session. When a new PDN is initiated, then existing PDN tracing stops and new PDN tracing starts. For it to happen, first the new PDN tracing is started and then the existing PDN tracing is stopped. Hence, new PSN and session manager sub-session ID is allocated.

For a multi-PDN call, when you start the MonSub with Multi-trace=ON, it traces the new PDN as part of new fastpath tracing session (that is MonSub session). Hence, after tracing the four PDN, MonSub CLI shows max tracing session reached. Tracing each PDN takes place as a separate MonSub session.

MonSub Statistics

A new mechanism is added to publish the statistics regarding the quality of FASTPATH PCAP capture on MonSub CLI. The new mechanism publishes the statistics whenever it receives the buffer full MEH indication at session manager, throttled at every five seconds. The feature supports a maximum of four buffers for a FASTPATH PCAP corresponding to MonSub session. The feature does not publish the statistics by default, and needs to be enabled through debug CLI on UPF.

- **debug uplane monsub-stats disabled**
- **debug uplane monsub-stats enabled**

The stats contains the following informations:

```
Packet accepted: 14250000          Packet rejected: 62297
Congestion Short Term: 0                    Congestion Longer Term: 0
Throttled: 0                                PCAP File Transfer Rate: 9.91 mbps
```

The PCAP file transfer rate is the rate at which copy Co-Proc writes the PCAP from RAM-FS to HD-RAID.

X-Header

This feature supports the X-Header capture in slowpath PCAP. When UPF inserts the X-Header for Uplink packet, the UPF captures the packet at entry and exit interfaces. So, the exit packet sent to N6 contains the inserted X-header.

When UPF inserts the X-header for Downlink packet, the UPF captures the packet at entry and exit interfaces. So, the exit packet sent to N3 contains the inserted x-header.

Configuration Procedure for Monitor Subscriber on UPF

The protocol monitor can be used to display information for a specific subscriber session that is currently being processed. Depending on the number of protocols monitored, and the number of sessions in progress, a significant amount of data is generated. It is highly recommended that logging be enabled on your terminal client in order to capture all of the information that is generated.

MonSub can also be initiated from UPF console. Monitoring for a given IMSI should not be enabled from both SMF and UP console.

Follow the instructions in this section to invoke and configure the protocol monitoring tool for a specific subscriber session.

Procedure

Step 1 Invoke the monitor subscriber command from the Exec mode by entering the **monitor subscriber** CLI command.

```
[local]host_name# monitor subscriber { callid | imei | imsi | ipaddr | ipv6addr |
msid | msisdn | next-call | pcf | peer-fa | peer-lac | sgsn-address | type |
username | smgr-instance }
```

An output listing all the currently available protocols, each with an assigned number, is displayed. Specify the method the monitor should use by entering the appropriate keyword.

Step 2 Specify the method the monitor should use by entering the appropriate keyword.

Select other options and/or enter the appropriate information for the selected keyword.

Step 3 Select other options and/or enter the appropriate information for the selected keyword.

If no session matching the specified criteria was being processed when the monitor was invoked, a screen of available monitoring options appears.

Step 4 Configure the amount of information that is displayed by the monitor. To enable or disable options, enter the letter or 2-digit number associated with that option (C, D, E, 11, 12, etc.). To increase or decrease the verbosity, use the plus (+) or minus (-) keys.

The current state, ON (enabled) or OFF (disabled), is shown to the right of each option.

Option **Y** for performing multi-call traces is only supported for use with the GGSN.

```
WARNING!!! You have selected options that can DISRUPT USER SERVICE
Existing CALLS MAY BE DROPPED and/or new CALLS MAY FAIL!!!
(Under heavy call load, some debugging output may not be displayed)
Proceed? - Select (Y)es or (N)o
```

Step 5 Repeat step 6 as needed to enable or disable multiple protocols.

Step 6 Press the **Enter** key to refresh the screen and begin monitoring.

The monitor remains active until disabled. To quit the protocol monitor and return to the prompt, press **q**.

Monsub CLI Options

Monitor Subscriber CLI – New Options

The following options with their default value are the additions to existing **monitor subscriber** command:

UPF Monitor Subscriber CLI

Following are the options:

- **W - UP PCAP Trace (ON)**: This parameter creates PCAP trace for slowpath and fastpath and is only applicable for CUPS.
- **U - Mon Display (ON)**: The non-protocol events (such as statistics, charging information from ECS, etc.) are also captured in slowpath PCAP files and are displayed on U-PLANE monitor console.

- **V - PCAP Hexdump (ON)**: This flag must be set to ON to capture the protocol packets in a text file in hexdump format on U-PLANE.

Note Currently, UP PCAP Trace flag must be set to ON to capture fastpath and slowpath PCAP files on CUPS.

Monitor Subscriber CLI – New Options

The following options with their default value are the additions to existing **monitor subscriber** command:

- **F - Packet Capture (Full Pkt)**: Captures all packets from fastpath.

  Using this option, operators can choose between full and partial packet captures. By entering **F**, the packet capture type can be changed to either full or partial. With partial packet capture, users can enter packet sizes from 1 to 16384 bytes. For example, if input is 20, only the first 20 bytes of fastpath packets will be captured and the remaining packets will be dropped.

Note When opening the PCAP file, the summary view displays full length of the packet, but the detailed view shows only the truncated packet.

- **/ - Priority (0)**: The value is in the range from "0 – Best Effort" to "7 – Guaranteed"
  - 0 - Best Effort
  - 1 - Low
  - 2 - Med-Low
  - 3 - Medium
  - 4 - Med-High
  - 5 - High
  - 6 - Critical
  - 7 - Guaranteed

Caution It is strongly recommended not to change the default value. It can adversely affect the system performance.

- **N - MEH Header (OFF):** The MEH header is stripped from the IP packet, when this option is configured

Show Monitor Subscriber Sessions

Following is the new CLI to show the ongoing MonSub session.

You can trigger the CLI **show monitor subscriber fastpath session all** from UPF.

- **SessId:** This is the local session id for MonSub session on UP session manager.
- **CallID:** Call id on Userplane.
- **PSN:** This is panopticon sequence number. There is a maximum of four MonSub fastpath tracing sessions on one UP with PSN ranging from 0-3.
- **Start time:** Time at which MonSub tracing session starts.
- **Interface Type:** This is to identify the call type, for which MonSub fastpath tracing session is started, whether it is Sxa, Sxb or Sxab.

Disconnect Monitor Subscriber Sessions

Following is the new CLI to disconnect the ongoing MonSub session. You can trigger the CLI from UPF.

monitor subscriber fastpath disconnect sessmgr-instance <UP SMGR Instance ID> session-id <Local Monitor subscriber Session ID at SMGR instance level>

If the MonSub session disconnect is successful, the following message dispalys on console.

```
Session Disconnected Successfully
```

If the MonSub session disconnect fails, the following message dispalys on console.

```
Monitor Subscriber session does not exist
```

Note Only security administrator can execute the monitor disconnect CLI.

Context, CDRMOD, and Hexdump Interaction for Monitor Subscriber

Hexdump module must be configured to provide operators the provision to configure Files names and Poll timers. The Hexdump module is one of the modules such as—EDR, UDR, and so on, that are part of the CDRMOD functionality. Configure the hexdump in a non-local context such as the ECS context. The local context does not support Hexdump modules.

For more information on Hexdump module and its configuration, refer to the Configuring the Hexdump Module for MonSub in UPF section.

PCAP File Name Convention

Following section discusses the naming conventions for PCAP files:

Note Only **monitor-subscriber-file-name** and **rotation** options are used in naming PCAP files.

Slowpath File Name Convention

The slowpath file names appear in the following format:

```
curr_slowpath_{SMGR Mon Sub Session
Id}_{monsub_file_name_option_val}_{Timestamp}_{RotationCount}.pcap
```

or

```
slowpath_{SMGR Mon Sub Session
Id}_{monsub_file_name_option_val}_{Timestamp}_{RotationCount}.pcap
```

File with ‘curr_’ prefix is the file, that is currently being written to, which is still not closed. When files are to be rotated (depending on the file rotation parameters), file without the ‘curr_’ prefix are copied to hard disk.

The SMGR MonSub Session Id – This is the session Id for MonSub session created on Uplane SMGR instance ID, which created this PCAP. This Id is local to SMGR instance, so there could be two SLOWPATH pcap captured with same ID.

When files are to be copied to hard disk, the monsub_file_name_option_val is replaced by:

- IMSI value if **monitor-subscriber-file-name** is set to "imsi".
- Call ID value if **monitor-subscriber-file-name** is set to "call-id"
- Username value if **monitor-subscriber-file-name** is set to ‘username’

Timestamp is in the following format "MMDDYYYYHHMMSS", where:

- MM - Month, DD - Date and YYYY - Year.
- HH -Hour, MM - Minutes and SS - Seconds.

RotationCount is a 9-digit value that is incremented every time an old file is rotated, and a new file is generated.

00000000 for the first file, 00000001 for the second file and so on.

Rotation of slowpath files is determined by following option in **hexdump-module file** configuration:

rotation { num-records *number* **| time** *seconds*| **volume** *bytes* **}**

- **num-records:** num-records specifies the number of packets after which, a new file is generated and ‘RotationCount’ in the filename is incremented. The range of number is between 100 to 10240, and the default value is 1024.
- **time:** time specifies the time to wait in seconds before a new file is generated and ‘RotationCount’ in the filename is incremented. seconds must be an integer from 30 through 86400. The default value is 3600.
- **volume:** volume specifies the number of bytes after which a new file is generated and ‘RotationCount’ in the filename is incremented. bytes must be an integer from 51200 through 62914560. The default value is 102400.



Note The **tarriff-time** parameter under rotation is ignored as it is not suitable for PCAP file capture.

The following are examples of the file naming conventions for slowpath PCAP files:

- For the ‘imsi’ option where IMSI is ‘112233445566778’, slowpath files are named as:

```
slowpath_S0_112233445566778_07152019050907_000000000.pcap
```

- For ‘call_id’ option where Call Id is ‘01317b22’, slowpath files are named as:

  `slowpath_S0_01317b22_07152019050907_000000000.pcap`



Note The parameter **tarrif-time** is not applicable for PCAP file capture.

Fastpath File Name Convention

The fastpath file names appear in the following format:

`vpp_{S}_{B}_{encap}_{monsub_file_name_option}_{Timestamp}_{FileCount}.pcap`

- S is replaced by either ‘S1’, ‘S2’,'S3’, or 'S4’.
- B is replaced by either ‘B0’,‘B1’,'B2', or 'B3' depending on the bundle generated by Panopticon.
- monsub_file_name_option is replaced by:
  - IMSI value if **monitor-subscriber-file-name** is set to "imsi".
  - Call ID value if **monitor-subscriber-file-name** is set to "call-id"
  - Username value if **monitor-subscriber-file-name** is set to ‘username’

Timestamp is in the following format "MMDDYYYYHHMMSS", where:

- MM - Month, DD - Date and YYYY - Year.
- HH -Hour, MM - Minutes and SS - Seconds.

RotationCount is a 9-digit value that is incremented every time an old file is rotated, and a new file is generated.

00000000 for the first file, 00000001 for the second file and so on.

Fast path "FileCount" is not the same as the slowpath "RotationCount" parameters and hence ‘hexdump-module file rotation’ parameters are ignored while naming fastpath files.

In Phases 1 of the feature, fastpath generated file names are like ‘vpp_S1_B0_ip.pcap’ or ‘vpp_S1_B1_ip.pcap’, they are renamed to following when being copied over to non-volatile storage:

- `vpp_S1_B0_ip_01317b22_07152019050907_000000000.pcap`
- `vpp_S1_B1_ip_01317b22_07152019050908_000000001.pcap`
- `vpp_S1_B0_ip_01317b22_07152019050908_000000002.pcap`

In MonSub phase 3, a PCAP “bundle” is replaced with a single PCAP file that uses Ethernet encapsulation.

In Phase 3, each fastpath session file is captured in the Ethernet PCAP file that is ‘vpp_S0_B0_eth.pcap’ and they are renamed to following when being copied to a non-volatile storage:

`vpp_S0_B0_eth_01317b22_07152019050907_000000000.pcap`

For ‘callid’ option where Call Id is ‘12345678ef’:

- `slowpath_S0_12345678ef _07152019050907_000000000.pcap`
- `vpp_S1_B0_eth_12345678ef _07152019050907_000000000.pcap`

For ‘username’ option where username is ‘9890098900’:

- `slowpath_S0_07152019050907_000000000_9890098900.pcap`
- `vpp_S1_B0_eth_07152019050907_000000000_9890098900.pcap`

PCAP File Location

Fastpath PCAP files are written to the `/records/pcap` directory in same card and CPU complex where the session manager owns the subscriber session resides.

`/records` directory is mapped to the "tmpfs" filesystem that is mapped to RAM. In this state, the files are suffixed with a ".pending" extension. For example:

```
-rw-rw-r-- 1 root root 268599296 Sep 23 14:04 vpp_S1_B0_eth.pending
```

Note The files size at this stage is not the actual file size when it is written to a persistent storage.

Once the fastpath tracing mechanism has written the files, they are converted to ‘.pcap’ files and renamed as given below. Additionally, there is a file that ends with a ".done" extension:

```
-rw-rw-r--    1 root     root       8689188 Oct 16 22:06 vpp_S0_B0_eth.pcap
```

After the PCAP files are written by fastpath tracing mechanism, the Co-Proc functionality instantiates and copies the files to a hard disk or a persistent storage.

The aforementioned file location process for Fastpath is also applicable to Slowpath.

The target file location in all cases is: `/hd_raid/records/hexdump`, except for the case in the hexdump module configuration where **use-harddisk** is enabled and the **directory** option under the **hexdump file** is to a custom value. For example, if the **directory** option is set to a value "abc" then the target location for the PCAP file will be: `/hd_raid/records/hexdump/abc/`.

In this feature implementation, a predefined location is set for PCAP files.

- To make sure that `/records/pcap` directory is not populated when issues are encountered with the use of **use-harddisk** and **hexdump module** configurations.
- For regular cleanup from `/hd_raid/records/hexdump` directory.

File Transfer to External Location

Once the files have been copied to the hard disk, they can be copied over to an external server using the command: **transfer-mode** option under the **hexdump** command in the **hexdump-module** configuration.

Apart from **transfer-mode**, other relevant options under **hexdump** can be used for external file transfer. Operators can use these commands to avoid excessive storage during fastpath processing.

Limitations

Some of the known limitations of the monitor subscriber feature are:

- This feature does not support SR/ICSR.

- Max of 4 mon sub sessions can be enabled per UPF.
- Monitoring for a given IMSI should not be enabled from both the SMF and the UP console.
- A combo call consists of 2 mon sub-sessions since it involves monitoring 2 PDNs (SXA and N4 legs with same IMSI but unique Call IDs).
- Incoming uplink packets (with GTPU) are not captured in VPP ingress in this phase.
- Restarting a trace immediately after quitting may result in fastpath files in `/records/pcap` the directory to be overwritten. Therefore, restart the session after a few seconds.
- When a MonSub trace stops, the tear down process may take a few seconds, so it is recommended to wait. It takes a maximum of (5 seconds, hexdump poll timer value in sec) before toggling the MonSub trace to start. Else, the operators may observe MAX TRACING SESSIONS REACHED momentarily.
- Monitor Subscriber fastpath sessions show CLI does not display the sessions that are being stopped. Hence, there is a transient period, where the new MonSub sessions can be rejected due to maximum sessions reached. However, show CLI shows fewer sessions. Therefore, wait for some time before starting a new MonSub trace session.
- Changing fastpath configuration options is only possible when **UP Pcap Trace** is set to OFF.
- When MT=ON in the Multi-PDN, then once MT=OFF, new PDN tracing is not started due to MAX TRACING REACHED, and then all other tracing is STOPPED. It happens because the first new PDN tracing starts and then all previous PDNs were STOPPED for MT=OFF case.
- Do not launch the same UE MonSub sessions from different CLIs.
- In slowpath PCAP, the egress DL packets do not show the GTP-U header, because the functionality to add GTP-U is available in fastpath. Hence, the ingress and egress DL packets show the duplicates, unless there are some packet modifications like HTTP X-headers applied over the ingress packets.
- Toggling C and D options does not impact the PCAP captures in UPF.
- For Multi-PDN, the fastpath file names do not use the Call ID. As by definition, the multi-PDN case has more than one call ID and hence a higher-level configuration such as IMSI is more suitable for naming the files.
- Only the named options that are explicitly mentioned in this document are supported from the *hexdump-module file* configuration.
- Number of streams that can be traced in fastpath is limited to 5000. A stream is defined as a TCP or UDP flow. It includes source IP address, destination IP address, source port, and destination port, transport protocol such as, TCP or UDP.
- Fastpath packets cannot be streamed to an external server. They are stored on the hard-disk and transferred either manually or by using **transfer-mode** options.
- The UP PCAP trace must be set to ON to capture fastpath and slowpath PCAP files.
- MonSub CLI option '<SPACE> Pause' is only to pause console events. There is no impact on other tracing events (slowpath PCAP, fastpath PCAP, and protocol packets tracing in a text file in hexdump format) with this option.
- The UP trace PCAP file does not contain the initial PFCP Sx Request/Response, due to race condition.

- The ICMP Packets and first packets of TCP and UDP streams flow through both slowpath and fastpath. Default values of GTPU (option 26) and User L3 (Option 19) are set to OFF. As a result, these packets are not captured in slowpath captures. If Option 26 is set to ON, then these packets are captured in slowpath PCAP captures. Option 19 has no effect on slowpath PCAP capture.
- Data Events flag must be set to ON, to capture fastpath and slowpath PCAP files.
- The MonSub tracing is not supported for option **Next-SAEGW Call** on UP.
- The MonSub tracing is not supported for option **Next call by APN** for Pure-S call type.
- On an ASR-5500 setup with the default value of poll-timer, all the packets may not be captured due to a known issue. To avoid rejection for a large number of packets, change the poll-timer value to the lowest possible (10 ms).
- If context replacement occurs (if the same subscriber reattaches without a detach), then the slowpath captures for the new call remain in the old slowpath files.

Configuring the Hexdump Module for MonSub in UPF

Configuring MonSub Poll Timer

Use this configuration to set the frequency of PCAP file capture check.

```
configure
  context context_name
    hexdump-module
      hexdump monitor-subscriber-poll-timeout poll_timer_value
      end
```

NOTES:

- **hexdump monitor-subscriber-poll-timeout** : This option specifies how frequently the check for newly captured PCAP files in the volatile storage must be done before they are copied to persistent storage.
- *poll_timer_value*: Specifies the poll timer value in milliseconds. It must be an integer in the range of 10 ms to 60 seconds. Default: 30 seconds.

Note The timer should not be configured with a value less than 5 seconds.

- This option is only applicable when MonSub is enabled for the products that have fastpath functionality - PGW, SAEGW on ASR-5500 and VPC-SI.

Configuring MonSub File Name

Use the following configuration to specify the file name of the PCAP file which contains IMSI, Call ID, or Username.

```
configure
  context context_name
    hexdump-module
      file rotation { num-records number | tariff-time minute minutes hour
 hours | time seconds| volume bytes | monitor-subscriber-file-name { imsi |
username | call-id }
       end
```

NOTES:

- **monitor-subscriber-file-name { imsi | username | call-id }**: This option specifies if the name of the captured PCAP files will contain IMSI, Call Id or Username. This option is only applicable on products that have fastpath functionality (PGW, SAEGW on ASR 5500 and VPC-SI) AND only when Monitor Subscriber functionality is enabled. Default: IMSI.
- **rotation { num-records** *number* **| tariff-time minute** *minutes* **hour** *hours* **| time** *seconds* **| volume** *bytes* **}**: Specifies when to close a hexdump file and create a new one.
  - **num-records** *number* : Specifies the maximum number of records that should be added to a hexdump file. When the number of records in the file reaches this value, the file is complete.

    number must be an integer from 100 through 10240. Default: 1024
  - **tariff-time minute** *minutes* **hour** *hours* : Specifies to close the current hexdump file and create a new one based on the tariff time (in minutes and hours).

    minutes must be an integer from 0 through 59.

    hours must be an integer from 0 through 23.
  - **time** *seconds* : Specifies the period of time to wait (in seconds) before closing the current hexdump file and creating a new one.

    seconds must be an integer from 30 through 86400. Default: 3600



Important Set the rotation time to 30 seconds.

  - **volume** *bytes* : Specifies the maximum size of the hexdump file (in bytes) before closing it and creating a new one.

    bytes must be an integer from 51200 through 62914560. Note that a higher setting may improve the compression ratio when the compression keyword is set to gzip. Default: 102400

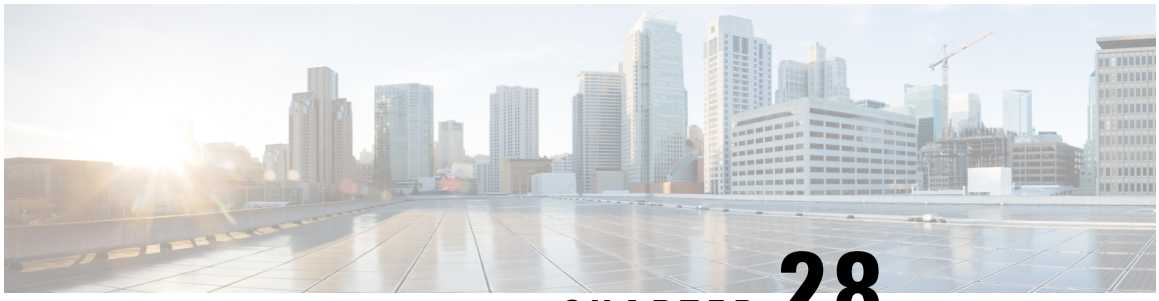
Monitoring and Troubleshooting

This section provides information regarding monitoring and troubleshooting the Monitor Subscriber feature.

SNMP Traps

The following SNMP trap(s) are added in support of the Monitor Subscriber feature:

- **MonSubProcessInitFailure:** This trap is triggered when MonSub handler process has failed for a particular process and service.



CHAPTER 28

MPLS Support on UPF

- Feature Summary and Revision History, on page 315
- Feature Description, on page 316
- How it Works, on page 316
- Monitoring and Troubleshooting, on page 323

Feature Summary and Revision History

Summary Data

Table 66: Summary Data

Applicable Product (s) or Functional Area	5G-UPF
Applicable Platforms	VPC-SI
Feature Default Setting	Disabled – Configuration Required
Related Changes in this Release	Not Applicable
Related Documentation	*UCC 5G UPF Configuration and Administration Guide*

Revision History

Revision Details	Release
Added MPLS support over the N9 interface.	2023.02.0
Added MPLS support over the N6 interface.	2023.01.0
First introduced.	2022.04.0

Feature Description

Multiprotocol Label Switching (MPLS) is a protocol that uses labels to route packets instead of using IP addresses. UPF supports MPLS to switch MPLS traffic using VPP as the data plane forwarder.

On ASR 5500, the NP4c network processor generates and processes MPLS traffic. On VPC-DI, IFTask generates and processes MPLS traffic.

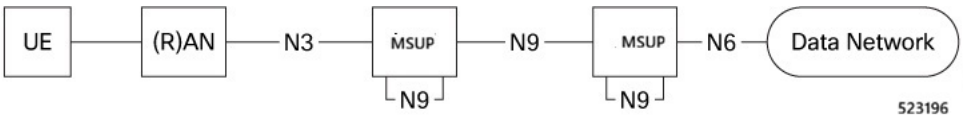
VPC-SI uses VPP as the data plane forwarder. VPP encapsulates and decapsulates subscriber traffic with MPLS labels. VPP provides support for multiple data plane features that include the MPLS stack as a separate graph node. This feature helps to differentiate between different customer VRFs and support corporate APNs having different addressing models and requirements.

UPF supports MPLS encapsulation over the following interfaces for 5G deployments using VPP:

- N6—N6 is the interface between UPF and data network (DN).
- N9—The N9 interface connects two UPFs. It is the interface between Intermediate I-UPF (Visited) and UPF Session Anchor (Home).

The following figure depicts the different interfaces used for MPLS encapsulation in UPF.

Figure 21: MPLS Interfaces



UPF supports the following functionalities for MPLS:

- Uses VPP MPLS stack to send the MPLS labeled packet.
- Uses VPP MPLS stack to process the incoming labeled MPLS packet.
- Uses only MP-BGP as the label distribution protocol.
- VPPCTL CLI commands display the FTN, FIB, and ILM tables that are in VPP, for debugging and comparing values with boxer configuration.
- In 1:1 redundancy deployment, UPF supports up to 129 VRFs on N6.
- In N:M redundancy deployment, UPF supports up to 129 VRFs on N6. The same 129 VRFs are present on each UP in the redundancy group.

How it Works

This section briefly describes how the MPLS feature works on UPF.

The VPP forwarder supports all existing functionalities for MPLS packet processing with its own MPLS stack. The VPP MPLS stack is configured with the appropriate FTN (FEC To NHLFE), forwarding information base (FIB) table, and incoming label map (ILM) table. The egress generates the MPLS packet with the correct MPLS header. It also processes the incoming MPLS packet and switches the packet to the appropriate VRF table based on the incoming label.

VPC-SI also supports VPNv6 as described in RFC 4659 – *BGP-MPLS IP Virtual Private Network (VPN) Extension for IPv6 VPN*.

Deploying UPFs using the N9 Interface

The 5G Core UPF can deploy two UPFs in series through the N9 interface. The N9 interface connects two UPFs between Intermediate I-UPF and UPF Session Anchor. To provide complete mobility with a stable IP anchor across the whole network, the two UPFs are connected depending on the operator network configuration.

Both upstream and downstream traffic between the UPFs is in the form of GTP-U packets. To support MPLS over the N9 interface, the L3 routers present between the two UPFs are capable of advertising BGP routes and labels to the UPFs.

In the upstream direction, I-UPF receives the GTP-U packets from gNodeB and modifies the content of the GTP-U header based on the destination (anchor) UPF. Based on the lookup of the IP header, MPLS encapsulation occurs to encapsulate the GTP-U packet. The MPLS label that needs to be inserted will be received through BGP from the intermediate L3 devices between UPFs. At the anchor UPF, the MPLS header gets decapsulated and transmitted as an IP packet towards the PDN.

In the downstream direction, MPLS encapsulation occurs in the anchor UPF after the IP packet received from PDN is encapsulated by a GTP-U header. This MPLS header is then decapsulated in I-UPF and the corresponding GTP-U packet is then transmitted to gNodeB.

In the following scenarios, UPF moves the N9 interface to a VRF and establishes a VPNv4 or VPNv6 AF BGP neighborship:

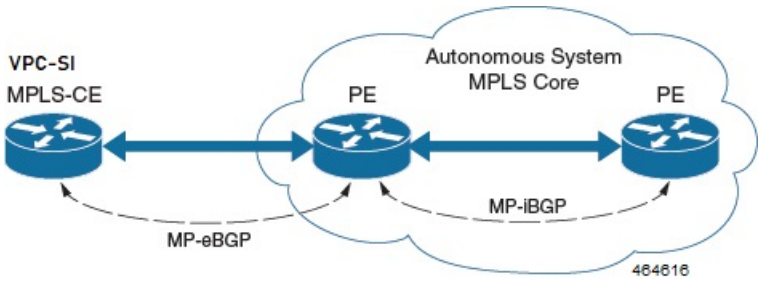
- When the N3 and N9 interfaces share the same physical interface.
- When the N3 and N9 interfaces are separated.

MPLS-CE Connected to PE

The VPC-SI functions as an MPLS-CE (Customer Edge) network element connected to a Provider Edge (PE) Label Edge Router (LER), which in turn connects to the MPLS core as per RFC 4364.

The following figure illustrates the MPLS-CE to PE connection:

Figure 22: VPC-SI MPLS-CE to PE



The MPLS-CE functions like a PE router within its own Autonomous System (AS). It maintains Virtual Routing and Forwarding (VRF) routes and exchanges VPN route information with the PE through an MP-eBGP (Multi Protocol external BGP) session.

The PE is also configured with VRFs and exchanges VPN routes with other PEs in its AS through MP-iBGP (Multi Protocol internal BGP) connection and MPLS-CE through an MP-eBGP connection.

The EBGP connection allows the PE to change next-hop IP addresses and labels in the routes learnt from IBGP peers before advertising them to the MPLS-CE. The MPLS-CE uses only MP-eBGP to advertise and learn routes. Label Distribution Protocol (LDP) and Resource Reservation Protocol (RSVP) are not required because of direct-connect EBGP peering. The MPLS-CE pushes or pops a single label (learnt over the MP-eBGP connection) to or from the PE.

Sample Configuration

MPLS N6 Configuration

The following is a sample configuration to configure the basic MPLS feature with N6 interface on UPF:

```
context egress
    ip vrf MPN00001
      ip maximum-routes 100
    #exit
mpls bgp forwarding
router bgp 64455
      neighbor 172.31.35.29 remote-as 65200
      neighbor 172.31.35.29 timers keepalive-interval 60 holdtime-interval 300
      neighbor 172.31.35.29 update-source 172.31.35.4
      neighbor 172.31.35.29 fall-over bfd
      neighbor 172.31.35.29 srp-activated-soft-clear
       address-family ipv4
        redistribute connected
        redistribute static
      #exit
      address-family vpnv4
        neighbor 172.31.35.29 activate
        neighbor 172.31.35.29 send-community extended
      #exit
      address-family vpnv6
        neighbor 172.31.35.29 activate
        neighbor 172.31.35.29 send-community extended
      #exit
      address-family ipv6
        redistribute connected
        redistribute static
      #exit
       ip vrf MPN00001
        route-distinguisher 64455 11100001
        route-target both 64455 11100001
      #exit
      address-family ipv4 vrf MPN00001
        redistribute connected
      #exit
      address-family ipv6 vrf MPN00001
        redistribute connected
      #exit
    interface N6-11
      ip address 172.31.35.4 255.255.255.224
      ip mtu 1504
      bfd interval 999 min_rx 999 multiplier 3
   #exit
    interface N6-12
      ip address 172.31.35.35 255.255.255.224
      ip mtu 1504
      bfd interval 999 min_rx 999 multiplier 3
    #exit
    interface loop_MPN00001 loopback
      ip vrf forwarding MPN00001
```

```
      ip address 2.2.4.2 255.255.255.255
    #exit
```

MPLS N9 Configuration

The following is a sample configuration to configure the MPLS feature with N9 interface on UPF:

```
context SAEGW
   ip vrf mpls-vrf-1
   #exit
   ip vrf mpls-vrf-2
   #exit
   mpls bgp forwarding
      bfd-protocol
      #exit
   ip prefix-list name up seq 15 permit 50.50.41.4/32
   ip prefix-list name up seq 20 permit 50.50.41.5/32
   ip prefix-list name up seq 25 permit 50.50.41.6/32
   ip prefix-list name up seq 40 permit 50.50.41.7/32
   ip prefix-list name n9-up1 seq 50 permit 60.60.41.5/32
   ip prefix-list name n9-up1 seq 55 permit 60.60.41.6/32
   route-map up_routes1 permit 10
      match ip address prefix-list n9-up1
      #exit
   route-map up_routes permit 10
      match ip address prefix-list up
      #exit
   ip vrf mpls-vrf-1
   route-distinguisher 61601 11100001
   route-target export 61601 11100001
   route-target import 61606 11100001
   route-target import 65200 11100001
   #exit
   address-family ipv4 vrf mpls-vrf-1
      redistribute connected
      redistribute static
      #exit
   address-family ipv6 vrf mpls-vrf-1
      redistribute connected
      redistribute static
      #exit
   ip vrf mpls-vrf-2
   route-distinguisher 61601 11100002
   route-target export 61601 11100002
   route-target import 61606 11100002
   route-target import 65200 11100002
   #exit
   address-family ipv4 vrf mpls-vrf-2
      redistribute connected
      redistribute static
      #exit
   address-family ipv6 vrf mpls-vrf-2
      redistribute connected
      redistribute static
      #exit
    #exit
   interface gn-mpls-n9-ingress
      ip address 172.31.40.5 255.255.255.224
      ip mtu 1504
      bfd interval 999 min_rx 999 multiplier 3
      #exit
   interface n9-ingress-hupf loopback
      ip vrf forwarding mpls-vrf-1
      ip address 60.60.41.5 255.255.255.255
```

```
   #exit
user-plane-service user_plane_svc
   associate gtpu-service pgw-gtpu upf-ingress
   associate gtpu-service n9-ingress upf-ingress interface-type n9
   associate gtpu-service sgw-gtpu-ingress-phazr sgw-ingress
   associate gtpu-service sgw-gtpu-egress sgw-egress
   associate gtpu-service SxC cp-tunnel
   associate sx-service sx-svc
   associate fast-path service
   associate control-plane-group SAEGW
   associate userplane-load-control-profile LCP
   associate userplane-overload-control-profile OLCP
   #exit
```

VPN-related CLI Commands

UPF supports the VPN-related features and functions across several CLI command modes. The following tables identify commands associated with configuration and monitoring of VPN-related functions.

Table 67: VPN-Related Configuration Commands

CLI Mode	Command	Description
BGP Address-Family (IPv4/IPv6) Configuration Mode	**neighbor** *ip_address* **activate**	Enables the exchange of routing information with a peer router.
BGP Address-Family (IPv4/IPv6) Configuration Mode	**neighbor** *ip_address* **send community { both | extended | standard }**	Sends the community attributes to a peer router (neighbor).
BGP Address-Family (IPv4/IPv6) Configuration Mode	**redistribute connected**	Redistributes routes into BGP from another protocol as BGP neighbors.
BGP Address-Family (VPNv4) Configuration Mode	**neighbor** *ip_address* **activate**	Enables the exchange of routing information with a peer router.
BGP Address-Family (VPNv4) Configuration Mode	**neighbor** *ip_address* **send community { both | extended | standard }**	Sends the extended-community attribute to a peer router. In VPN, route-distinguisher and route-target are encoded in the BGP extended-community. This command sends BGP routes with the extended community to a neighbor.
BGP Address-Family (VRF) Configuration Mode	**neighbor** *ip_address* **activate**	Enables the exchange of routing information with a peer router.
BGP Address-Family (VRF) Configuration Mode	**neighbor** *ip_address* **send community { both | extended | standard }**	Sends the extended-community attribute to a peer router. In VPN, route-distinguisher and route-target are encoded in the BGP extended-community. This command sends BGP routes with the extended community to a neighbor.

CLI Mode	Command	Description
BGP Address-Family (VRF) Configuration Mode	**redistribute connected**	Redistributes routes into BGP from another protocol as BGP neighbors.
BGP Configuration Mode	**address-family { ipv4 vrf** *vrf_name* **| vpnv4 }**	Enables the exchange of IPv4 VRF routing information. There is a different mode for each address-family.
BGP Configuration Mode	**address-family { ipv6 vrf** *vrf_name* **| vpnv6 }**	Configures a VPNv6 address family and IPv6 VRF routing in BGP.
BGP Configuration Mode	**ip vrf** *vrf_name*	Adds a VRF to BGP and switches to the VRF Configuration mode to allow configuration of BGP attributes for the VRF.
BGP IP VRF Configuration Mode	**route-distinguisher {** *as_value* **|** *ip_address* **}** *rd_value*	Assigns a Route Distinguisher (RD) for the VRF. The RD value must be a unique value on the router for each VRF.
BGP IP VRF Configuration Mode	**route-target { both | import | export } {** *as_value* **|** *ip_address* **}** *rd_value*	Adds a list of import and export route-target extended communities to the VRF.
Context Configuration Mode	**ip pool** *pool_name addr_range* **vrf** *vrf_name* **[mpls-label input** *inlabel1* **output** *outlabel1 outlabel2* **]**	Configures a pool into the specified VRF. Specify this command with the Next-Hop parameter. *inlabel1* is the MPLS label that identifies inbound traffic destined for this pool. *outlabel1* and *outlabel2* specify the MPLS labels to be added to packets sent for subscribers from this pool.
Context Configuration Mode	**ip vrf** *vrf_name*	Creates a VRF and assigns a VRF-ID. A VRF is created in the router.
Context Configuration Mode	**ipv6 pool** *pool_name* **vrf** *vrf_name*	Associates the pool with that VRF. **Note:** By default the configured IPv6 pool is associated with the global routing domain.
Context Configuration Mode	**mpls bgp forwarding**	Globally enables MPLS Border Gateway Protocol (BGP) forwarding.

CLI Mode	Command	Description
Context Configuration Mode	**mpls exp** *value*	Sets the default behavior as Best Effort using a zero value in the 3-bit MPLS EXP header. This value applies to all the VRFs in the context. The default behavior is to copy the DSCP value of mobile subscriber traffic to the EXP header, if there is no explicit configuration for DSCP to EXP (via the **mpls map-dscp-to-exp dscp** *n* **exp** *m* command). **mpls exp** disables the default behavior and sets the EXP value to the configured *value*.
Context Configuration Mode	**mpls ip**	Globally enables the MPLS forwarding of IPv4 packets along normally routed paths.
Context Configuration Mode	**radius change-authorize-nas-ip** *ip_address* **{ encrypted | key }** *value* **port** *port_num* **mpls input** *inlabel* **output** *outlabel1 outlabel2*	Configures COA traffic to use the specified MPLS labels. *inlabel* identifies the inbound COA traffic. *outlabel1* and *outlabel2* specify the MPLS labels to add to the COA response. *outlabel1* is the inner output label and *outlabel2* is the outer output label.
Ethernet Interface Configuration Mode	**mpls ip**	Enables dynamic MPLS forwarding of IP packets on this interface.
Exec Mode	**clear ip bgp peer**	Clears BGP sessions.

Table 68: VPN-Related Monitoring Commands

CLI Mode	Command	Description
Exec Mode show Commands	**show ip bgp neighbors**	Displays information regarding BGP neighbors.
Exec Mode show Commands	**show ip bgp vpnv4 { all | route-distinguisher | vrf }**	Displays all VPNv4 routing data, routing data for a VRF or a route-distinguisher.
Exec Mode show Commands	**show ip bgp vpnv6**	Displays contents of VPNv6 routing table.
Exec Mode show Commands	**show ip bgp vpnv6 { all | route-distinguisher | vrf }**	Displays all VPNv6 routing data, routing data for a VRF or a route-distinguisher.
Exec Mode show Commands	**show ip fib table-id** *table_id*	Displays information for FIB table.

CLI Mode	Command	Description
Exec Mode show Commands	**show ip pool**	Displays pool details including the configured VRF.
Exec Mode show Commands	**show mpls cross-connect**	Displays MPLS cross-connect information. MPLS tunnel cross-connects between interfaces and Label-Switched Paths (LSPs). LSP connects two distant interface circuits of the same type through MPLS tunnels.
Exec Mode show Commands	**show mpls fib table** *table_id*	Displays MPLS forwarding information base (FIB) table information.
Exec Mode show Commands	**show mpls ftn [vrf** *vrf_name* **]**	Displays MPLS FEC-to-NHLFE (FTN) table information.
Exec Mode show Commands	**show mpls ftn [vrf** *vrf_name* **]**	Displays the contents of the MPLS FTN table for a specified VRF.
Exec Mode show Commands	**show mpls ilm [all | label** *label_value* **]**	Displays MPLS Incoming Label Map (ILM) table information.
Exec Mode show Commands	**show mpls nexthop-label-forwarding-entry**	Displays MPLS Next-Hop Label Forwarding Entry (NHLFE) table information.

Monitoring and Troubleshooting

This section provides information regarding the CLI command to monitor and troubleshoot the feature.

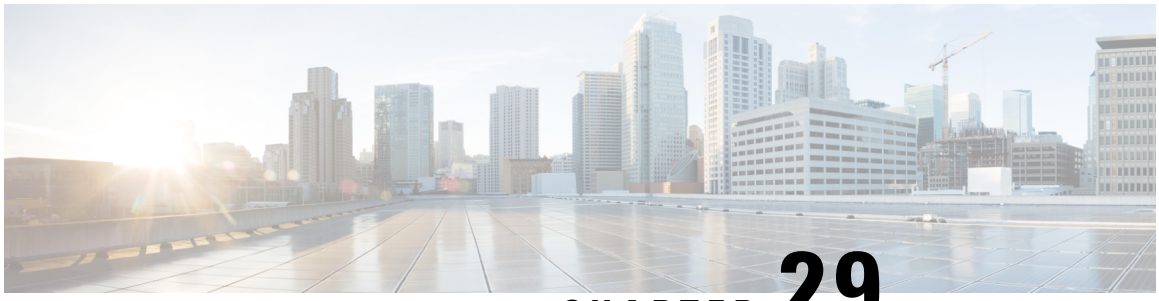
Show Commands and Outputs

This section provides information regarding show commands and their outputs in support of this feature.

show mpls ftn vpp

The output of this command contains the "vpp" field that displays the configured VPP dataplane values in the VPP dataplane forwarder. Use this command for debugging purposes.

- vpp
  - all-vrf
  - summary
  - vrf



CHAPTER 29

Multiple cnSGW Support

- Feature Summary and Revision History, on page 325
- Feature Description, on page 326
- How it Works, on page 326
- Configuring Multiple SMF/cnSGWs, on page 326
- Monitoring and Troubleshooting, on page 327

Feature Summary and Revision History

Summary Data

Table 69: Summary Data

Applicable Product (s) or Functional Area	5G-UPF
Applicable Platforms	VPC-SI SMI
Feature Default Setting	Disabled – Configuration Required
Related Changes in this Release	Not Applicable
Related Documentation	*UCC 5G UPF Configuration and Administration Guide*

Revision History

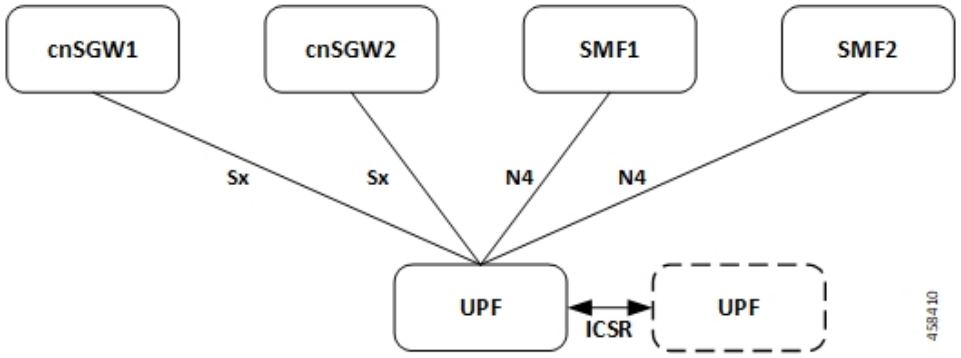
Revision Details	Release
First introduced.	2021.02.0

Feature Description

The Multiple SMF/cnSGW Support feature enables a single UPF to establish multiple N4/Sx interfaces with cnSGWs and their paired SMFs. Integration of multiple SMF and cnSGW combinations with a single UPF results in optimal usage of resources.

Architecture

The following illustration depicts the architecture of multiple cnSGWs/SMFs.



Relationship to Other Features

The Multiple cnSGW Support feature is related to *Multiple N4/Sx Interface* feature.

How it Works

The functionality of Multiple cnSGW feature involves:

- Single UPF has multiple N4/Sx interface associations with SMF/cnSGWs.
- There's no slicing of configuration in UPF per individual SMF.
- Cumulatively, a maximum of four peers—combination of cnSGW/SMF or individual cnSGW/SMF as per the need—are connected to a single UPF.
- Individual N4/Sx association release purges sessions of the impacted peer.
- UPF redundancy works seamlessly.
- All cnSGWs paired with a UPF is associated with a single user plane service.

Configuring Multiple SMF/cnSGWs

This section provides information about CLI commands that are available in support of this feature.

Configuring Multiple SMF/cnSGWs on UPF

Use the following CLI commands to configure multiple SMF/cnSGWs on UPF by adding multiple peer node under Control Plane Group Configuration mode.

```
configure
  user-plane-service service_name
    associate control-plane-group group_name
  control-plane-group group_name
    peer-node-id ipv4-address ipv4_address interface n4
    peer-node-id ipv4-address ipv4_address interface n4
    peer-node-id ipv4-address ipv4_address
    peer-node-id ipv4-address ipv4_address
    end
```

Monitoring and Troubleshooting

This section provides information about monitoring and troubleshooting the Multiple cnSGW feature.

Show Commands and/or Outputs

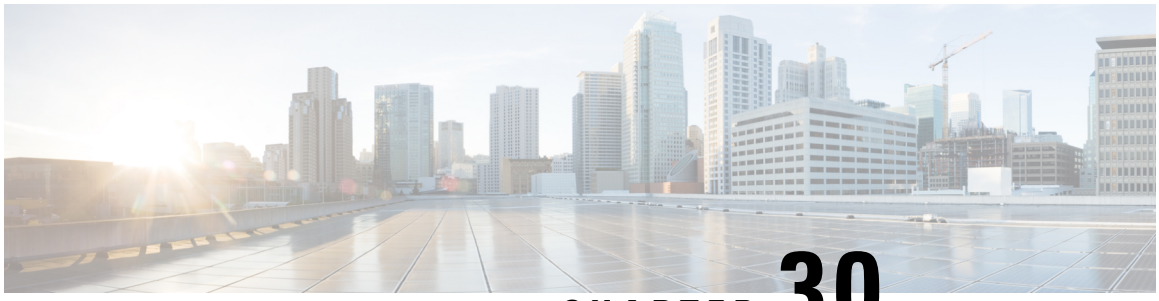
This section describes the show commands that are available in support of this feature.

show subscribers user-plane-only full all

The output of this CLI command is enhanced to display the corresponding Control Plane address.

show sx peers

The output of this CLI command is enhanced to display the peer ID with corresponding number of sessions.



CHAPTER 30

Multiple N4/Sx Interfaces

- Feature Summary and Revision History, on page 329
- Feature Description, on page 330
- How it Works, on page 330
- Configuring Multiple N4/Sx Interfaces, on page 331
- Monitoring and Troubleshooting, on page 332

Feature Summary and Revision History

Summary Data

Table 70: Summary Data

Applicable Product (s) or Functional Area	5G-UPF
Applicable Platforms	VPC-SI SMI
Feature Default Setting	Disabled – Configuration Required
Related Changes in this Release	Not Applicable
Related Documentation	*UCC 5G UPF Configuration and Administration Guide*

Revision History

Revision Details	Release
UPF supports the following functionality: • Multiple N4/Sx interfaces with any number of control plane NFs • Maximum configuration of 18 Nx/Sx peer nodes	2023.04.0

Revision Details	Release
First introduced.	2021.01.0

Feature Description

Table 71: Feature History

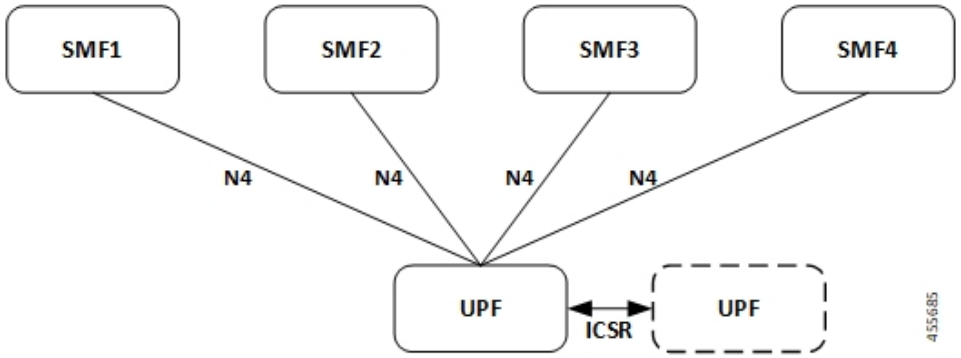
Feature Name	Release Information	Description
Support for Multiple N4/Sx Interfaces	2023.04	A single UPF can establish multiple N4 or Sx interfaces with any number of control plane network functions such as SMF, cnSGWc, SAEGW-C, PGW-C, and SGW-C. The maximum number of supported N4/Sx peer nodes has been increased from 16 nodes to 18 nodes in this release.

The Multiple N4/Sx Interfaces feature enables a single UPF to establish multiple N4 or Sx interfaces with any number of control plane NFs. The control plane NFs include SMF, SAEGW-C, PGW-C, SGW-C, and cnSGWc.

The integration of multiple control plane nodes with a single UPF results in optimal usage of resources. UPF supports a maximum configuration of 18 N4 or Sx peer nodes per control-plane-group.

Architecture

The following figure illustrates the architecture of multiple N4 interfaces with a single UPF.



How it Works

The functionality of the Multiple N4/Sx Interfaces feature involves:

- The ECS/ACS configuration at UPF is a union of all individual control plane specific configurations.

For example:

SMF1 has rulebase *RB1* and no *RB2*

SMF2 has rulebase *RB2* and no *RB1*

The UPF has both rulebase, *RB1* and *RB2* to cater the sessions from *RB1* and *RB2*.

- There is no slicing of configuration in UPF per individual SMF.
- A maximum number of four SMF peers are connected to a single UPF.
- Overlapping IP pools from multiple SMFs are segregated based on the VRF ID.
- Individual N4 association release purges sessions of the impacted SMF peer.
- UPF redundancy works seamlessly.
- During any conflict with different SMF configurations, it will be installed in the sequence of configured CLIs and not resolved at UPF.

Configuring Multiple N4/Sx Interfaces

This section provides information about CLI commands that are available in support of this feature.

Configuring Multiple Peer Nodes

Use the following CLI commands to configure multiple SMF on UPF by adding multiple peer node under Control Plane Group Configuration mode.

To configure multiple peer nodes on UPF, use the following sample configuration:

```
configure
  user-plane-service service_name
    associate control-plane-group group_name
  control-plane-group group_name
    peer-node-id ipv4-address ipv4_address interface n4
    peer-node-id ipv4-address ipv4_address interface n4
    peer-node-id ipv4-address ipv4_address
    . . .
    . . .
    . . .
    end
```

NOTES:

- **peer-node-id ipv4-address** *ipv4_address*—Specify the IPv4 address of the peer node.
- **interface n4**—Identify the N4 interface.
- A maximum number of 18 peer nodes can be configured per control plane group.

  When the limit exceeds beyond 18 peer nodes, he following error message displays:

  Failure: Maximum Control Plane Group Nodes Limit exceeded!

Monitoring and Troubleshooting

This section provides information about monitoring and troubleshooting the Multiple N4/Sx Interface feature.

Show Commands and/or Outputs

This section describes the show commands that are available in support of this feature.

show ip chunks

The output of this CLI command is enhanced to display the IP pools pushed to the UPF from multiple SMFs in Gi context.

show ipv6 chunks

The output of this CLI command is enhanced to display the IPv6 pools pushed to the UPF from multiple SMFs in Gi context.

show subscribers user-plane-only full all

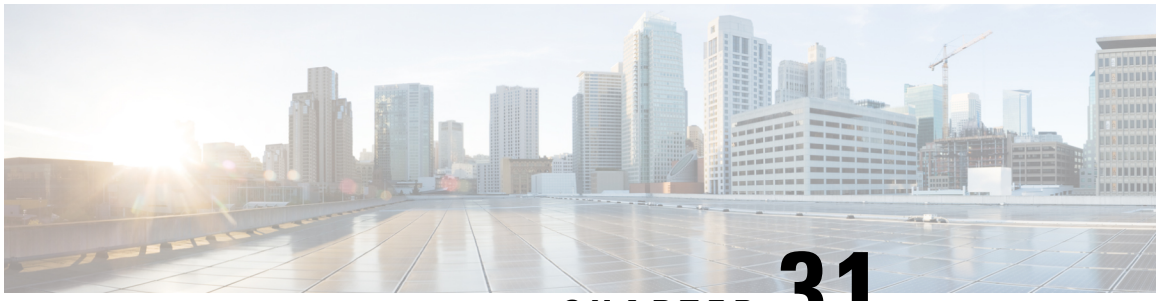
The output of this CLI command is enhanced to display the corresponding Control Plane address.

show sx peers

The output of this CLI command is enhanced to display the peer ID with corresponding number of sessions.

show user-plane-service statistics peer-address <address>

The output of this CLI command is enhanced to display per peer statistics in SMF.



CHAPTER 31

Nexthop Forwarding Support

- Revision History, on page 333
- Feature Description, on page 333
- How It Works, on page 333
- Configuring Next Hop Forwarding Support, on page 338
- VLAN Segmentation Using Next Hop Forwarding, on page 342
- Monitoring and Troubleshooting, on page 346

Revision History

Table 72: Revision History

Revision Details	Release
First introduced.	2022.01.0

Feature Description

The next hop forwarding is the process of forwarding the packets to an adjacent router or device. The next hop router is the next closest or optimal IP for a packet in the destination path.

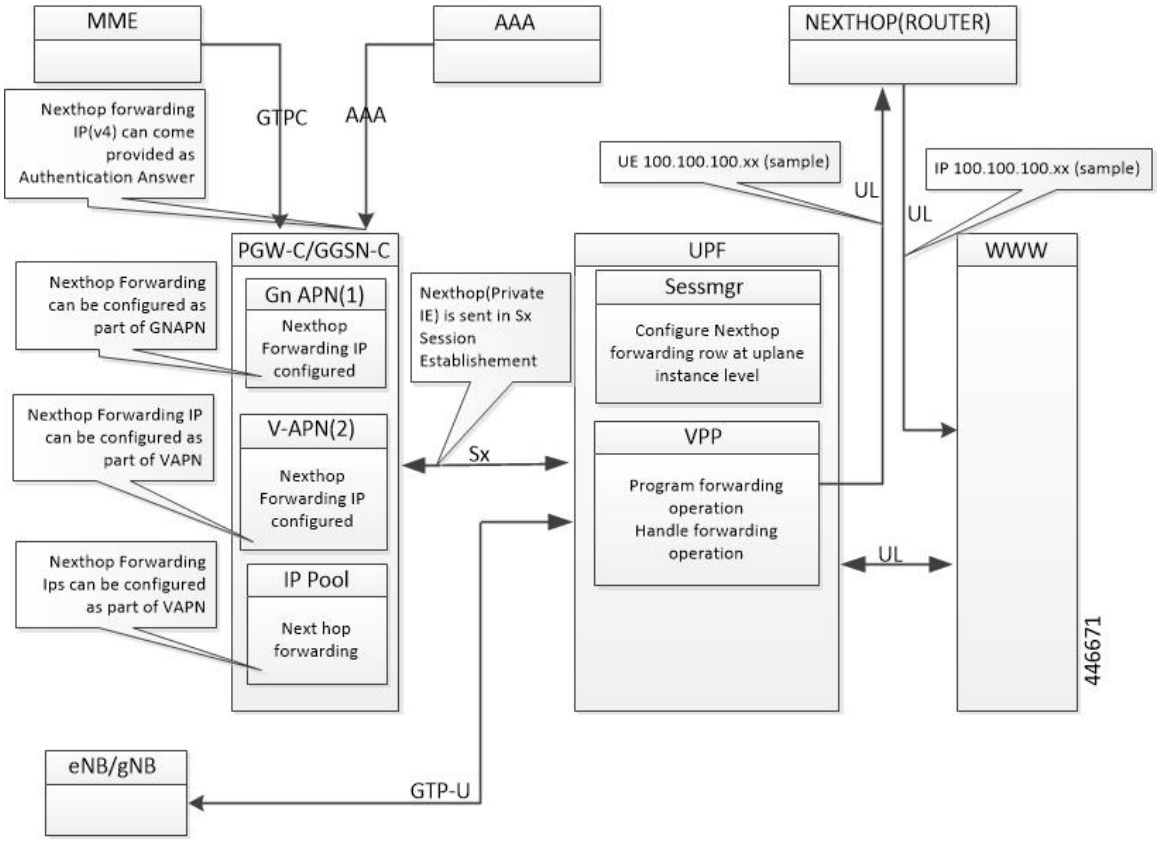
In the uplink direction, the UE and the GI IP can be in a different subnet. The routing path in UPF is defined to allow the uplink packet to be forwarded accordingly.

How It Works

Architecture

The following illustration provides an overview of Next hop forwarding feature.

Figure 23: Next hop forwarding



You can configure Next hop IP address at the SMF under DNN profile or IPAM profile. During PDU Establishment, the SMF relays the IPv4/IPv6 address over the N4 interface in the NEXT HOP IP private IE in a PFCP Session Establishment Request.

You can also configure the Next hop IP address at the UPF through Charging-Action.

Note When Next hop address is provided by both SMF and UPF Charging-Action, the UPF Charging-Action Nexthop address takes precedence.

Configuration Priority

S. No.	Configuration	Priority
1.	UPF (Charging Action)	1
2.	DNN profile	2
3.	IP Pool	3

Configuration Use Cases

Case	IP Type	DNN	IP Pool	UPF (Charging Action)	Nexthop IP Selection
Next hop supplied only in DNN	IPv4	209.165.201.18	Not configured	Not configured	Next hop Address is selected from DNN: IPv4: 209.165.201.18 IPv6: 8001::10
	IPv6	8001::10	Not configured		
Next hop supplied only in IP pool	IPv4	Not configured	209.165.201.19	Not configured	Next hop Address is selected from IP Pool: IPv4: 209.165.201.19 IPv6: 8001::10
	IPv6	Not configured	8001::10		
IPv4 and IPv6 configured in DNN and IP pool	IPv4	209.165.201.18	209.165.201.19	Not configured	Next hop Address is selected from DNN: IPv4: 209.165.201.18 IPv6: 9001::3
	IPv6	9001::3	8001::10		
IPv6 configured in DNN and IPv4 in IP pool	IPv4	Not configured	209.165.201.19	Not configured	Next hop IPv4 is selected from IP pool: 209.165.201.19 Next hop IPv6 selected from DNN : 8001::10
	IPv6	8001::10	Not configured		
IPv6 configured in DNN, IPv4 in IP pool, and IPv4 in UPF Charging Action	IPv4	209.165.201.18	209.165.201.19	209.165.201.20	Next hop Address is selected from UPF (CA): 209.165.201.20
	IPv6	8001::10	Not configured	Not configured	
IPv6 configured in DNN, IPv4 in IP pool, and IPv6 in UPF Charging Action	IPv4	209.165.201.18	209.165.201.19	Not configured	Next hop Address is selected from UPF(CA) : 9001::10
	IPv6	8001::10	Not configured	9001::10	

Interface

The following Private IEs are introduced in Sx/N4 Session Establishment message.

238	PFCP_IE_NEXTHOP	PFCP_IE_NEXTHOP								Sx/N4 Session Establishment Request	Private IE: UPF: nexthop forwarding support-IPv4 /IPv6 address
		BITS									
		Octets	7	6	5	4	3	2	1		
		1 to 2	Type = 238 (decimal)								
		3 to 4	Length = n								
		5 to 10	PFCP_IE_NEXTHOP_ID								
		11-14	PFCP_IE_NEXTHOP_IP								

239	PFCP_IE_NEXTHOP_ID	PFCP_IE_NEXTHOP_ID								1. Inside Create FAR IE of Sx Session Establishment Request	Private IE : UPF: nexthop forwarding support-IPv4 /IPv6 address
		BITS								2. Inside PFCP_IE_NEXTHOP IE of Sx/N4 Session Establishment Request	
		Octets	7	6	5	4	3	2	1	PFCP_IE_NEXTHOP of Sx/N4 Session Establishment Request	
		1 to 2	Type = 239 (decimal)								
		3 to 4	Length = 5								
		5 to 10									

240	PFCP_IE_NEXTHOP_IP	PFCP_IE_NEXTHOP_IP									
		Bits								PFCP_IE_NEXTHOP of Sx/N4 Session Establishment Request	Private IE : UPF: nexthop forwarding support-IPv4/ IPv6 address
		Octets	7	6	5	4	3	2	1		
		1 to 2	Type = 240 (decimal)								
		3 to 4	Length = n								
		5	spare					V4	V6		
		m to m+3	IPv4 Address								
		p to p+15	IPv6 Address								

The following is a sample output of SX_SESSION_ESTABLISHMENT_REQUEST in which the SMF relays the Next hop IP to the UPF.

```
CREATE FAR:
  Type: 3
  Value:
   FAR ID:
     Type: 108
     Value: 0x80000002
   APPLY ACTION:
     Type: 44
     Value:
       DROP: 0
       FORW: 1
       BUFF: 0
       NOCP: 0
       DUPL: 0
   FORWARDING PARAMETERS:
     Type: 4
     Value:
       DESTINATION INTERFACE:
         Type: 42
         Value: CORE (1)
       PDN INSTANCE:
          Type: 22
          Value: intershat
          INNER PACKET MARKING:
            Type: 220
            TOS/TRAFFIC CLASS: 0xB8 0xFC
    NextHop ID:
       Value: 0x0001

CREATE TRAFFIC ENDPOINT:
    Type: 127  Length: 20
    Value:
    Traffic Endpoint ID:
        Type: 131  Length: 1
        Value: 0x0004
        Hex: 0083 0001 04
```

```
        Local F-TEID:
            Type: 21  Length: 1
            Value:
                CH: 1
                IPv4: 0
                IPv6: 0
                CHID: 0
            Hex: 0015 0001 04
        Bearer Info:
            Type: 225  Length: 6
            QCI: 5
            ARP: 84
            Charging ID: 5592407
            Hex: 00E1 0006 0554 0055 5557
    Hex: 007F 0014 0083 0001 0400 1500 0104 00E1
                 0006 0554 0055 5557
NEXT HOP IP:
    Type: 237  Length: 14
    Value:
    NextHop ID:
        Type: 238  Length: 1
        Value: 0x0001
        Hex: 00EE 0001 01
    IP ADDR:
        Type: 239  Length: 5
        Value:
            IPv4: 1
            IPv6: 0
            IPv4: 209.165.202.150
            IPv6:
            Hex: 00EF 0005 020F 0F0F 0F
    Hex: 00ED 000E 00EE 0001 0100 EF00 0502 0F0F
         0F0F
```

Limitations

The following are the known limitations of this feature:

- Configuring next hop forwarding through AAA is not supported.
- Next hop address sent on RADIUS and Diameter (Redirect information from PCF) interfaces is not qualified.
- When you configure next hop forwarding in DNN profile and IPAM, the next hop is only seen in Sx Establishment, and not in Create FAR IE of Sx Session Modification Request.

Configuring Next Hop Forwarding Support

Configuring Next Hop Forwarding Through Charging Action

At the UPF, use the following CLI commands to configure Next hop forwarding through Charging Action.

```
configure
  active-charging service service_name
    charging-action charging_action_name
      nexthop-forwarding-address ipv4_address/ipv6_address
```

NOTES:

- **charging-action** *charging_action_name*: Specifies the name of a charging action. *charging_action_name* must be an alphanumeric string of 1–63 characters and can contain punctuation characters. Each charging action must have a unique name.
- **nexthop-forwarding-address** *ipv4_address/ipv6_address*: Configures the next hop forwarding address.

Configuring Next Hop Forwarding Through DNN Profile

At the SMF, use the following CLI commands to configure Next hop forwarding through the DNN profile.

```
configure
  profile dnn intershat
       nexthop-forwarding-address { ipv4 ipv4_address| ipv6  ipv6_address }
       end
```

NOTES:

- **nexthop-forwarding-address { ipv4** *ipv4_address* **| ipv6** *ipv6_address* **}**: Configures the Next hop forwarding address.
  - *ipv4_address*: Configures IPv4 address.
  - *ipv6_address*: Configures IPv6 address (supports colon-separated hexadecimal notation).

Configuring Next Hop Forwarding at IP Pool Through IPAM Profile for IPv4 Addresses

At the SMF, use the following CLI command to configure Next hop forwarding at the IP pool through the IPAM profile for IPv4 addresses.

```
configure
  ipam
    instance instance_id
      address-pool pool_name
        ipv4
        address-ranges
          address-range start_ipv4_address end_ipv4_address
nexthop-forwarding-address nexthop_forwarding_address
          prefix-range prefix_value length  prefix_length
nexthop-forwarding-address nexthop_forwarding_address
          split-size per-cache number_of_addresses
          split-size per-dp number_of_addresses
          commit
```

NOTES:

- **address-pool** *pool_name*: Specify the name of the address pool. *pool_name* must be a string.
- **ipv4**: Enter the IPv4 mode of the pool.
- **address-ranges**: Specify the starting address of the IPv4 address range. Enter the IPv4 address range and prefix range addresses with the next hop forwarding address.
  - **address-range** *start_ipv4_address end_ipv4_address* **nexthop-forwarding-address** *nexthop_forwarding_address*: Specify the starting and the ending addresses of the IPv4 address range with the next hop forwarding address.
  - **prefix-range** *prefix_value* **length** *prefix_length*: Specify the prefix value and the length within the IPv4 address.
  - **nexthop-forwarding-address** *nexthop_forwarding_address*: Specify the next hop forwarding address.
- **split-size per-cache** *number_of_addresses*: Specify the number of IPv4 addresses per chunk for IPAM cache allocation. Specify in the power of 2. The IPAM server consumes this configuration. *number_of_addresses* must be an integer in the range of 2-262144.
- **split-size-per-dp** *number_of_addresses*: Specify the number of IPv4 addresses per chunk for data plane allocation. Specify in the power of 2. The IPAM cache consumes this configuration.

  number_of_addresses must be an integer in the range of 2-262144.

Configuring Next Hop Forwarding at IP Pool Through IPAM Profile for IPv6 Addresses

To configure the IPv6 address with the next hop configuration for IPv6 pools and address ranges, use the following sample configuration:

```
configure
   ipam
     instance instance_id
        address-pool pool_name
           ipv6
           address-ranges
              address-range start_ipv6_address end_ipv6_address
nexthop-forwarding-address nexthop_forwarding_address
              prefix-range prefix_value length prefix_length
nexthop-forwarding-address nexthop_forwarding_address
                split-size per-cache number_of_addresses
                split-size per-dp number_of_addresses
                exit
                prefix-range prefix_value length prefix_length
nexthop-forwarding-address nexthop_forwarding_address
              commit
```

NOTES:

- **address-pool** *pool_name*: Specify the name of the address pool. *pool_name* must be a string.
- **ipv6**: Enter the IPv6 mode of the pool.

- **address-ranges**: Specify the IPv6 address ranges and prefix range addresses with the next hop forwarding address.

Note IANA IPv6 configuration is used by BNG.

  - **address-range** *start_ipv6_address end_ipv6_address* : Specify the starting and the ending addresses of the IPv6 address range.
  - **nexthop-forwarding-address** *nexthop_forwarding_address*: Specify the next hop forwarding address.
  - **prefix-range** *prefix_value* **length** *prefix_length*: Specify the prefix value and length within the IPv6 address.
  - **nexthop-forwarding-address** *nexthop_forwarding_address*: Specify the next hop forwarding address.

- **prefix-ranges** : Specify the prefix ranges of an IPv6 address.

Note SMF supports only IAPD IPv6 configuration.

  - **split-size per-cache** *number_of_addresses*: Specify the number of IPv6 addresses per chunk for IPAM cache allocation.
  - **split-size-per-dp** *number_of_addresses*: Specify the number of IPv6 addresses per chunk for the data plane allocation.
  - **prefix-range** *prefix_value* **length** *prefix_length* **nexthop-forwarding-address** *nexthop_forwardng_address*: Specify the prefix value and the length within the IPv6 address with the next hop forwarding address.

Configuration Example

The following is an example configuration to configure next hop forwarding at IP pool through IPAM profile for IPv4 addresses:

```
config
   ipam
      instance 1
         address-pool p1
                  ipv4
                     split-size per-cache 1024
                     split-size per-dp 256
                     end
```

Following is an example configuration to configure next hop forwarding at IP pool through IPAM profile for IPv6 addresses:

```
ipam
 instance 1
  address-pool ISE-Pool1
   vrf-name ISP
```

```
tags
 dnn cisco_vlan400.com
exit
ipv6
 address-ranges
   address-range 1000::1 1000::ffff nexthop-forwarding-address :9001::3
   prefix-range 2607:fc20:1010:: length 98 nexthop-forwarding-address :9001::3
 prefix-ranges
    split-size
        per-cache 32768
        per-dp    32768
    exit
    prefix-range 2607:fc20:1010:: length 44 nexthop-forwarding-address :9001::3
 exit
exit
```

VLAN Segmentation Using Next Hop Forwarding

Table 73: Feature History

Feature Name	Release Information	Description
VLAN Segmentation of Traffic Using Next Hop Forwarding	2024.01	UPF allows VLAN segmentation using Next Hop Forwarding. UPF sends packets to the VLAN from where the next hop IP is most reachable instead of sending packets to the destination IP. **Default Setting**: Disabled-Configuration required to enable

Feature Description

UPF supports segmenting the enterprise traffic using Virtual LAN (VLAN). VLANs enable a single physical LAN segment to be further segmented into smaller virtual networks so that the groups of ports are isolated from one another, as if they were on physically different segments. In VLAN segmentation, the UEs belonging to different classes are assigned to different VLANs for traffic isolation.

VLAN Segmentation can be achieved using two methods:

1. Next Hop Forwarding
2. Virtual Routing and Forwarding

In the VLAN segmentation using next hop forwarding, the IP pool is mapped with the VLAN tags. This mapping takes place at the SMF. The IP pool is associated with the next hop forwarding address on the N6 interface, which is reachable through the VLAN associated with the IP pool as configured on the Control Center.

Note This feature is supported on both 1-Rack Unit and 3-Rack Unit architectures and can support up to 16 VLANs on the N6 interface.

How it Works

Achieving VLAN segmentation using next hop forwarding follows the given process:

- The SMF sends a N4 Session Establishment Request with the next hop forwarding address to the UPF.
- The UPF associates the next hop forwarding address with the session. This association allows the session management layer on the UPF to program the VPP to perform next hop forwarding of the uplink packets belonging to that session.
- The session manager on the UPF creates the uplink streams in the VPP Fastpath conduit with next hop forwarding operation with appropriate next hop IP. Therefore, the UPF does not look at the actual IP destination route to forward the packets. Rather, it relates the UE's IP address with the next hop IP and forwards the uplink packet on the VLAN from where the next hop IP is reachable.

Enabling VLAN Segmentation Using Next Hop Forwarding

To configure VLAN segmentation on UPF using next hop forwarding, use the following configuration:

```
config
  port ethernet slot/port
    no shutdown
    vlan vlan_tag_ID
    no shutdown
    bind interface interface_name context_name
    end
    context context_name
      { ip | ipv6 } address address subnetmask [ secondary ]
      end
```

NOTES:

- **port ethernet** *slot*/*port*—Ethernet port or slot that will contain the VLAN tag.
- **no shutdown**—It enables or disables the traffic over a specified VLAN.
- **vlan** *vlan_tag_ID*—Enters the VLAN configuration mode.
- **no shutdown**—Enables or disables traffic over the current VLAN.
- **bind interface** *interface_name context_name*—It binds a virtual interface and context to support VLAN service.
- **context** *context_name*—Configures the context to in which the interface is created.
- **{ ip | ipv6 } address** *address subnetmask* **[secondary]** —Maps the IP pools with the configured VLANs. The **[secondary]** keyword assigns multiple IP addresses to the interface.

Configuration Example

To enable the VLAN segmentation using next hop forwarding on UPF, use the following sample configuration:

```
config
     port ethernet 1/10
       no shutdown
```

```
       vlan 400
         no shutdown
         bind interface N6_interface1 ISP
       #exit
       vlan 401
         no shutdown
         bind interface N6_interface2 ISP
       #exit

context ISP
 interface N6_interface1
      ip address 209.165.200.225 209.165.200.254
      ip address 209.165.201.1 209.165.200.254 secondary
      ipv6 address 2001:DB8::1/32 secondary

interface N6_interface2
      ip address 209.165.202.129 209.165.200.254
      ip address 209.165.201.30 209.165.200.254 secondary
      ipv6 address 2001:DB8:1::1/32 secondary
```

Configuring Network Reachability Server on UPF

In order to forward the packets to the next hop IP, it requires the UPF to perform source-based routing and also learn the MAC address of the next hop IP by sending an ARP request to the next hop.

Configuring network reachability allows the UPF to test the network device by sending the ping packets to the destination.

To test and detect network reachability or failure on UPF for IPv4 addresses, use following configuration:

```
config
  context context_name
  nw-reachability server server_name [ remote-addr ip_addr ] [ local-addr
 ip_addr ]
```

NOTES:

- **nw-reachability server** *server_name* **[remote-addr** *ip_addr* **] [local-addr** *ip_addr* **]** —This CLI adds or deletes a reachability-detect server. The **[local-addr** *ip_addr* specifies the source IP from which the packets are sent and the **[remote-addr** *ip_addr* specifies the destination IP to send the ping packets for detecting network failure or reachability.

Note UPF supports a maximum of 2000 IPv4 pools and 256 IPv6 pools per VPN network reachability configurations per context.

To test and detect network reachability or failure on UPF for IPv6 addresses, use the following configuration:

```
config
  context context_name
    ip access-list acl_name
    { deny | permit } [ log ] any
    exit
  nw-reachability server server_name [ remote-addr ip_addr ] [ local-addr
 ip_addr ]
  end if_name
```

```
interface if_name
  ipv6 address ipv6_address ipv6_mask
    bfd interval interval_value min_rx rx_value multiplier multiplier_value
    exit
  port ethernet slot/port
    no shutdown
    vlan vlan_tag_ID
    no shutdown
    bind interface interface_name context_name
    end
```

NOTES:

- **ip access-list** *acl_name*: **ip access-list** specifies the type of subscriber traffic and the direction (uplink, downlink, or both) traffic is redirected. The **ip access-list** must be specified in the context in which the subscriber authentication is performed.
- **{ deny | permit } [log] any**: Indicates the rule, when matched, drops or allows the corresponding packets. **[log]** indicates that all the packets which match the filter are to be logged.

  any indicates that all the packets will match the filter regardless of the source and/or the destination.
- **interface** *if_name*: Specifies the interface in which the to enable BFD.
- **ipv6 address** *ipv6_address ipv6_mask*: Specifies the IPv6 address for configuring the static routes.
- **bfd interval** *interval_value* **min_rx** *rx_value* **multiplier** *multiplier_value*: **bfd interval** *interval_value* specifies the transmit interval of control packets in milliseconds as an integer from 50 through 10000. **min_rx** *rx_value* specifies the minimum receive interval for control packets in milliseconds as an integer from 50 through 10000. **multiplier** *multiplier_value* specifies the value used to compute hold-down time as an integer from 3 through 50.

Configuration Example

Following is the sample configuration of network reachability to learn the MAC address of the next hop IP:

```
(config-ctx)# nw-reachability server test remote-addr 209.165.200.225 local-addr
209.165.202.129
```

OAM Support

This section describes the operations, administration, and maintenance support for this feature.

Monitoring Support

This section describes feature-level monitoring support information.

show ip arp vpp

This show command displays the correct MAC address associated with the next hop IP addresses:

```
# show ip arp vpp
Wednesday November 22 13:03:40 EST 2023
    Time        IP      Flags     Ethernet             Interface                  slot/port
3460.6250  10.1.105.254  D   00:fd:22:53:01:81  VirtualFunctionEthernet0/6/0.1105   1/10
3446.2276  10.2.105.254  D   00:fd:22:52:ff:5f  VirtualFunctionEthernet0/7/0.2105   1/11
```

show configuration context

To verify the configured ACL lists, use the following show command:

```
[ISP]Hertz-UPF-A# show configuration context  ISP
  context ISP
    ip access-list ecs-v4.acl
      permit any
    #exit
    nw-reachability server NHFWDVlan1 remote-addr 192.11.55.100 local-addr 192.11.55.51
    nw-reachability server NHFWDVlan1_ipv6 remote-addr 2001:4888:192:1155::100 local-addr
2001:4888:192:1155::51
    #exit
  interface isp-1155
      ip address 192.11.55.51 255.255.255.0
      ipv6 address 2001:4888:192:1155::51/64 secondary
      bfd interval 50 min_rx 50 multiplier 3

  port ethernet 1/10
    no shutdown
    vlan 1155
      no shutdown
      bind interface isp-1155 ISP
    #exit
```

show nw-reachability server

This show command displays the reachability status of both the next hop IPs (remote-addr). If the next hop IP is reachable, it is displayed as **Up.**

```
[ISP]upf01#  show nw-reachability server all
 Server               remote-addr             local-addr              Vrf           state
 ---------------      ---------------     ---------------- ---------------- ---------------
 vlan2645             10.28.45.1              10.28.45.2              n/a             Up
 vlan2646             10.28.46.1              10.28.46.2              n/a             Up
NHFWDVlan1_ipv6  2001:4888:192:1155::100 2001:4888:192:1155::51      n/a             Up

Total Network Reachability Servers: 3 Up: 3
```

show ipv6 interface summary

This show command displays the interface summary for VLAN configuration with IPv6 configuration:

```
[ISP]Hertz-UPF-A# show ipv6 interface summary
Interface Name               Address/Mask               Port               Status
====================    ========================    ==================    ============
=========
isp                     2001:4888:192:1153::51/64      1/10 vlan 1153          UP
 [sec]
isp-1155                2001:4888:192:1155::51/64      1/10 vlan 1155          UP
 [sec]
```

Monitoring and Troubleshooting

This section provides information about CLI commands available for monitoring and troubleshooting the feature.

Show Commands and Outputs

This section provides information about show commands and their outputs in support of this feature.

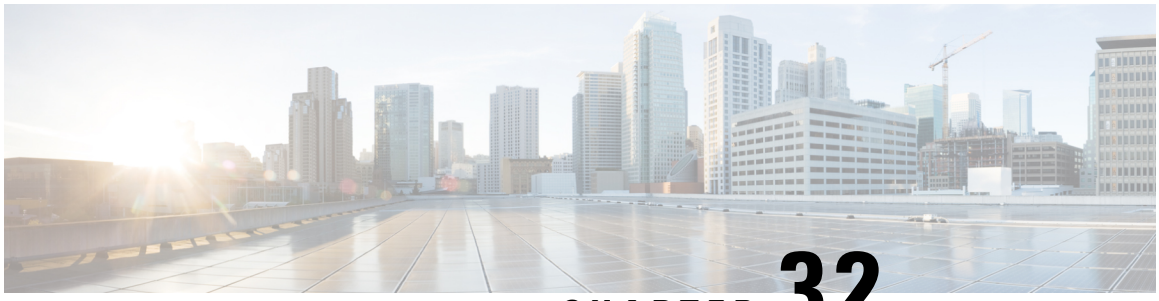
show subscriber user-plane-only full all

The output of this show command is enhanced to include the following fields introduced in support of this feature.

- **Next Hop Ip Address** - Displays the configured Nexthop IP address.



Note **Next Hop Ip Address** field is displayed only if the Nexthop IP address is relayed from the SMF. This field is not displayed if Nexthop IP address is configured only at the UPF using Charging-Action.



CHAPTER 32

N:M Redundancy and Redundancy Configuration Manager

- Feature Summary and Revision History, on page 349
- Feature Description, on page 350

Feature Summary and Revision History

Summary Data

Table 74: Summary Data

Applicable Product(s) or Functional Area	5G-UPF
Applicable Platform(s)	VPC-SI
Feature Default Setting	Disabled – Configuration Required
Related Changes in this Release	Not Applicable
Related Documentation	*UCC 5G RCM Configuration and Administration Guide*

Revision History

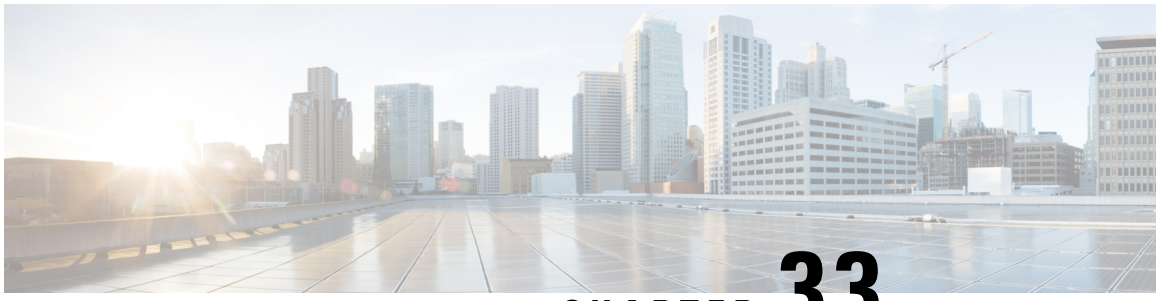
Table 75: Revision History

Revision Details	Release
Support has been added to defer SSH IP installation.	2021.02.2
Support has been added for configuring the RCM through Network Services Orchestrator (NSO).	2021.02.0
First introduced.	2021.01.0

Feature Description

The Redundancy Configuration Manager (RCM) is a Cisco proprietary node/network function (NF) that provides redundancy of StarOS-based UP/UPFs. The RCM provides N:M redundancy of UP/UPFs wherein “N” is the number of Active UPs/UPFs and is less than 10, and “M” is the number of Standby UP/UPF in the redundancy group.

For details, refer the Redundancy Configuration Manager - Configuration and Administration Guide.



CHAPTER 33

N3 Transfer of PDU Session Information

- Feature Summary and Revision History, on page 351
- Feature Description, on page 351

Feature Summary and Revision History

Summary Data

Applicable Product(s) or Functional Area	5G-UPF
Applicable Platform(s)	VPC-SI SMI
Feature Default Setting	Enabled – Always-on
Related Changes in this Release	Not Applicable
Related Documentation	Not Applicable

Revision History

Revision Details	Release
First introduced.	2020.02.0

Feature Description

The N3 transfer of PDU session information involves the inclusion of QoS Field Identifier (QFI) IE in the GTP-U extension header while performing GTP-U encapsulation toward gNodeB on the N3 interface, and removal of the GTP-U extension header while performing GTP-U decapsulation when packets are received from the gNodeB.

The QFI IE detects traffic pertaining to specific QoS sessions. It is used to send control information between the gNodeB and the UPF.

How it Works

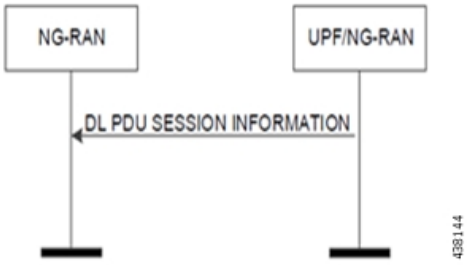
This section describes the transfer of PDU session Information procedures between the gNodeB and UPF for Uplink and Downlink packets.

Transfer of PDU Session Information for Downlink Data Packets

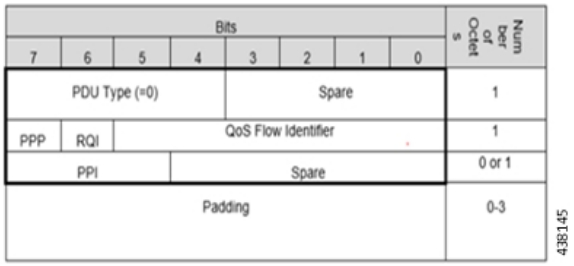
The Transfer of PDU Session Information for downlink data packets involves transfer of control information elements related to the PDU Session from UPF/NG-RAN to NG-RAN.

A PDU session user plane instance that makes use of this transfer procedure is associated to a single PDU Session. The procedure is invoked whenever packets for that particular PDU Session must be transferred across the related interface instance.

The DL PDU SESSION INFORMATION frame includes a QoS Flow Identifier (QFI) field that is associated with the transferred packet. The NG-RAN uses the received QFI to determine the QoS flow and QoS profile which are associated with the received packet.



The following frame shows the respective DL PDU SESSION INFORMATION.



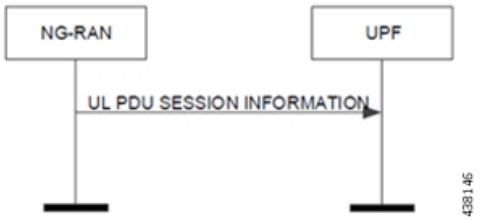
NOTE: In current implementation, the Reflective QoS Indicator (RQI) and Paging Policy Presence (PPP) in DL PDU SESSION INFORMATION frame is not supported.

Transfer of PDU Session Information for Uplink Data Packets

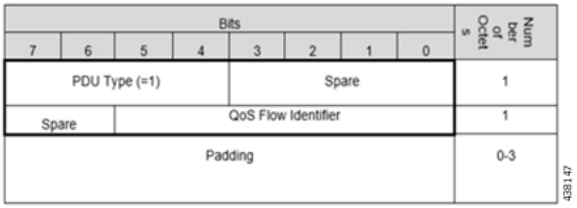
The Transfer of PDU Session Information for uplink data packets involves transfer of control information elements related to the PDU Session from NG-RAN to UPF.

An UL PDU Session user plane instance that makes use of the transfer procedure is associated to a single PDU Session. This procedure is invoked whenever packets for that particular PDU Session need to be transferred across the related interface instance.

The UL PDU SESSION INFORMATION frame includes a QoS Flow Identifier (QFI) field associated with the transferred packet.



The following frame shows the respective UL PDU SESSION INFORMATION.



PDU Session Information Frame IEs

The following table describes the Information Elements present in the PDU Session Information frame.

Information Element	Description
PDU Type	The PDU Type indicates the structure of the PDU session UP frame. The field takes the value of the PDU Type it identifies: "0" for PDU Type 0. The PDU type is in bit 4 to bit 7 in the first octet of the frame. **Value range**: {0= DL PDU SESSION INFORMATION, 1=UL PDU SESSION INFORMATION, 2-15=reserved for future PDU type extensions} **Field length**: 4 bits
Spare	The spare field is set to "0" by the sender and should not be interpreted by the receiver. This field is reserved for later versions. **Value Range**: (0–2n-1) **Field Length**: n bits
QoS Flow Identifier	When this IE is present, this parameter indicates the QoS Flow Identifier of the QoS flow to which the transferred packet belongs. **Value range**: {0 to 2^6-1} **Field length**: 6 bits
Padding	The padding is included at the end of the frame to ensure that the PDU Session user plane protocol PDU length (including padding and the future extension) is (n*4– 2) octets, where n is a positive integer. **Field Length**: 0–3 octets.

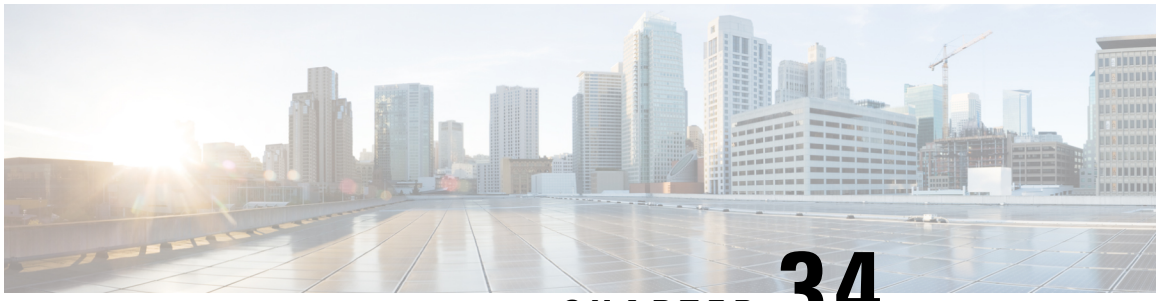
Standards Compliance

The feature complies with the following standard: 3GPP TS 38.415 V15.2.0 (NG-RAN; PDU Session User Plane Protocol).

Limitations

The following are the known limitations to this feature in this release:

- Reflective QoS Indicator (RQI) is not supported in this release.



CHAPTER 34

N4 Interface Compliance with 3GPP Specification

This chapter covers the following topics:

- Feature Summary and Revision History, on page 355
- Feature Description, on page 356

Feature Summary and Revision History

Summary Data

Table 76: Summary Data

Applicable Product(s) or Functional Area	5G-UPF
Applicable Platform(s)	VPC-SI SMI
Feature Default Setting	Enabled - Always-on
Related Changes in this Release	Not Applicable
Related Documentation	Not Applicable

Revision History

Table 77: Revision History

Revision Details	Release
Support is added for Outer Header Removal IE.	2021.04.0
In this release, PFCP library is upgraded to support the latest version of Outer Header IE.	2020.02.5
First introduced.	2020.02.0

Feature Description

In compliance with 3GPP TS 29.244, the User Plane Function (UPF) supports the following IEs:

- Averaging Window
- Paging Policy Indicator (PPI)
- Outer Header Creation
- Outer Header Removal

Averaging Window

Averaging window IE contains the duration over which the GBR and MBR is calculated. It is sent from SMF to UPF with Create QER or Update QER parent IE, if the default pre-configured value under UPF needs to be overridden.

The following format is used for encoding and decoding of the IE:

	Bits							
Octets	8	7	6	5	4	3	2	1
1 to 2	Type = 157 (decimal)							
3 to 4	Length = n							
5 to 8	Averaging Window							
9 to (n+4)	These octet(s) is/are present only if explicitly specified							

NOTE: The value should be in milliseconds.

Paging Policy Indicator

The SMF sends PPI value in Create QER or Update QER, if UPF requires to set Paging Policy Indicator in outgoing packets.

In the case of Network Triggered Service Request and UPF buffering downlink data packet, the UPF includes the DSCP in ToS (IPv4) / TC (IPv6) value from the IP header of the downlink data packet. It also sends an indication of the corresponding QoS Flow in the data notification message to the SMF. When PPD applies, the SMF determines the Paging Policy Indicator (PPI) based on the DSCP received from the UPF.

In the case of Network Triggered Service Request and SMF buffering downlink data packet, when PPD applies, the SMF determines the PPI based on the DSCP in ToS (IPv4) / TC (IPv6) value from the IP header of the received downlink data packet and identifies the corresponding QoS Flow from the QFI of the received downlink data packet.

The following format is used for encoding and decoding of the IE:

	Bits							
Octets	8	7	6	5	4	3	2	1

1 to 2	Type = 158 (decimal)	
3 to 4	Length = n	
5	Spare	PPI value
6 to (n+4)	These octet(s) is/are present only if explicitly specified	

NOTE: The PPI should be encoded as a 3-bit value from 0 through 7.

Outer Header Creation

Per 3GPP TS 29.244 v16.4.0, the Outer Header Creation Description field, when present, is encoded as specified in following table. It takes the form of a bitmask where each bit indicates the outer header to be created in the outgoing packet. Spare bits are ignored by the receiver.

Octet / Bit	Outer Header Created in the Outgoing Packet
5/1	GTP-U/UDP/IPv4
5/2	GTP-U/UDP/IPv6
5/3	UDP/IPv4
5/4	UDP/IPv6
5/5	IPv4
5/6	IPv6
5/7	C-TAG
5/8	S-TAG
6/1	N19 Indication
6/2	N6 Indication
6/3	TCP/IPv4
6/4	TCP/IPv6

NOTE:

- Currently, the UP/UPF doesn't support the following values of Outer Header Creation Description:
  - IPv4
  - IPv6
  - C-TAG
  - S-TAG
  - N19 Indication
  - N6 Indication
- Third and fourth bits of sixth Octet (that is, 6/3 and 6/4) are spare bits (that is, not part of 3GPP TS) used for LI over TCP.

- When the Outer Header Creation IE is received with the Outer Header Creation description field set to both GTP-U/UDP/IPv4 and GTP-U/UDP/IPv6, then the length of the Outer Header Creation IE, which is 26 bytes, is validated. If the length is not 26 bytes, then the Sx message is responded with the PFCP_CAUSE_MANDATORY_IE_INCORRECT with Offending IE: OUTER_HDR_CREATION cause code.



Important If SMF/CP uses older version for Outer Header Creation, then undefined behavior (including crashes) can be seen.

Outer Header Removal

Outer Header Removal feature is used to remove the GPRS Tunneling Protocol User Plane (GTP-U) header from the uplink GTP-U packets.

The following format is used for encoding Outer Header Removal Information Element (IE):

	Bits							
Octets	8	7	6	5	4	3	2	1
1–2	Type = 95 (decimal)							
3–4	Length = n							
5	Outer Header Removal Description							
6	GTP-U Extension Header Deletion							
7 to (n+4)	These octets are present only if explicitly specified							

Per 3GPP TS 29.244, the Outer Header Removal Description field, when present, is encoded as specified in the following table.

Table 78: Outer Header Removal Description

Outer Header to be Removed from the Incoming Packet	**Value (Decimal)**
GTP-U/UDP/IPv4 (See Notes 1, 2),	0
GTP-U/UDP/IPv6 (See Notes 1, 2)	1
UDP/IPv4 (See Notes 3, 6)	2
UDP/IPv6 (See Notes 3, 6)	3
IPv4 (See Note 6)	4
IPv6 (See Note 6)	5
GTP-U/UDP/IP (See Note 4)	**6**
VLAN S-TAG (See Note 5)	7
S-TAG and C-TAG (See Note 5)	8

Outer Header to be Removed from the Incoming Packet	Value (Decimal)
For future use. Not sent. If received, it's interpreted as the value "1".	9–255

NOTES:

- The SGW-U/I-UPF stores GTP-U extension headers. These headers are forwarded for the packets that aren't requested to be deleted by the GTP-U Extension Header Deletion field.
- The SGW-U/I-UPF stores the GTP-U message type for a GTP-U signaling message, which must be forwarded. For example, an End Marker message
- This value applies to DL packets received by a PGW-U for non-IP PDN connections. These connections use SGi tunneling based on UDP/IP encapsulation.
- The CP function uses this value for instructing the UP function to remove the GTP-U/UDP/IP header regardless of the IP version (IPv4 or IPv6).
- This value applies to DL packets received by a UPF over N6 for Ethernet PDU sessions.
- This value applies to DL packets received by a UPF (PDU Session Anchor) over N6, when explicit N6 traffic routing information is provided to the SMF.
- The Local F-TEID in the Created PDR or Traffic Endpoint IE is populated with the local GTP-U endpoint. This endpoint is allocated based on the configuration irrespective of the Outer Header Removal received in the Created PDR IE. For example, if the GTP-U service configuration supports both the IPv4 and IPv6 addresses, then both these addresses are allocated and advertised in the Created PDR and Traffic Endpoint IEs.

  In addition to the Local F-TEID allocation process, the User Plane advertises the allocated F-TEID to the Control Plane through the Created PDR or Traffic Endpoint IE with the following scenarios:

  - The Outer Header Removal that is received in the Create PDR IE is also considered while populating the Created PDR or Traffic Endpoint IE.
  - If Outer Header Removal is received as GTP-U, UDP, or IP, then the local GTP-U endpoint is advertised as allocated, which implies that if both the IPv4 and IPv6 addresses are configured in the GTP-U service, then both these addresses are allocated and advertised. If an IPv4 or IPv6 address is configured, then only the required address is allocated and advertised through the Created PDR or Traffic Endpoint IE.
  - If Outer Header Removal is received as GTP-U, UDP, or IPv4, then only the IPv4 address is advertised even if both these addresses are configured.
  - If Outer Header Removal is received as GTP-U, UDP, or IPv6, then only the IPv6 address is advertised even if both these addresses are configured.

Software Requirements

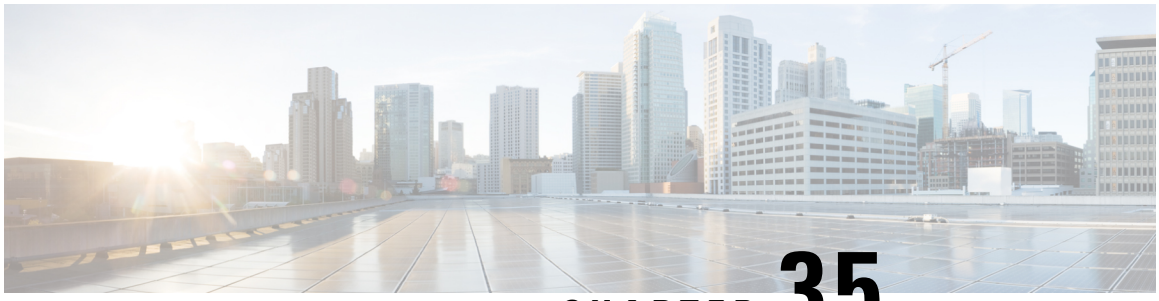
The software requirements are as follows:

- The feature requires UPF support to identify, encode, and decode the wildcard tunnel type "GTP-U/UDP/IP-6" on the N4 interface.
- If IPv4 and IPv6 addresses are received as part of Outer Header Creation (OHC), priority is given to the IPv6 endpoint, and hence the IPv6 Outer Header Removal (OHR) endpoint is retained by the UPF.

- GTP-U/UDP/IP-6 on the N4 interface, can be received over Sx Establishment or Sx Modification request messages. UPF must support type-6 on both cases.
- In Handoff scenarios, for all the PDRs with OHR value- 6, uplink packets are buffered until an appropriate OHC IE is received for PDRs corresponding to the downlink FAR.
- The uplink packets are forwarded only after the appropriate OHR type is set at UPF.

Limitations

- When the Outer Header Removal value as 6 is received for uplink PDR, the UPF maintains only IPv6 Outer Header Removal IE for uplink PDR. The UPF maintains it until an appropriate Outer Header Creation IE is received for downlink FAR.
- This feature is applicable to the N4 interface only.



CHAPTER 35

N4 Interface Configuration

This chapter covers the following topics:

- Feature Summary and Revision History, on page 361
- Feature Description, on page 362
- Uniformity in compression at N4 and Sx interfaces, on page 362
- Configuring N4 Interface, on page 363

Feature Summary and Revision History

Summary Data

Table 79: Summary Data

Applicable Product(s) or Functional Area	5G-UPF
Applicable Platform(s)	VPC-SI SMI
Feature Default Setting	Enabled (configuration required to disable)
Related Changes in this Release	Not Applicable
Related Documentation	*UCC 5G UPF Configuration and Administration Guide*

Revision History

Table 80: Revision History

Revision Details	Release
New IEs are supported in UPF in compliance with 3GPP TS 29.244.	2021.01.0
First introduced.	2020.02.0

Feature Description

This chapter provides the configuration information to identify a peer node to be an N4 interface, and the configuration to modify N4 parameters in an Sx-Service.

Uniformity in compression at N4 and Sx interfaces

This feature introduces a peer-level handshake mechanism during the PFCP Association Setup procedure. This allows the UPF and the Control Plane (CP) to negotiate compression capabilities, enabling the UPF to send compressed or uncompressed messages on a per-peer basis.

In deployments where a User Plane Function (UPF) is shared between a Converged Core (cnSGW/SMF) and a legacy Control and User Plane Separation (CUPS) core, call rejections may occur if Sx compression is enabled on the UPF. This is because converged core nodes do not currently support compressed PFCP messages.

How it works

The feature relies on an exchange of Information Elements (IEs) during the PFCP Association Setup procedure:

1. UPF Capability: The UPF sends the Operator configurable UPF capability IE (Type 359, as defined in 3GPP TS 29.244) in the Association Setup Request.
2. CP Capability: The CP (SMF/cnSGW) responds with a proprietary Operator configurable CP capability IE (Type 32769, Enterprise ID 9) in the Association Setup Response.
3. Capability Logic:
   - If the CP reports "0", the UPF sends uncompressed messages to that peer.
   - If the CP reports "1" or the IE is absent (backward compatibility), the UPF sends compressed messages, provided the local CLI is configured to allow it.

Per-Peer compression logic

The UPF determines whether to compress outgoing PFCP messages based on the following logic:

Local CLI config	Peer capability (IE 32769)	UPF compression
Disabled	Not Sent / 0 / 1	Disabled
Not Configured	Not Sent / 1	Enabled
Not Configured	0	Disabled
Enabled	Not Sent / 1	Enabled
Enabled	0	Disabled

Configuration

To enable or disable Sx protocol compression on the UPF, use the following CLI commands:

```
configure
          context <context_name>
```

```
sx-service <service_name>
  [no] sx-protocol compression
   end
```



Note The `sx-protocol compression` CLI controls outgoing message compression only. The compression flag (IE 359/IE 32769) indicates the node's capability to receive compressed messages. Runtime configuration changes take effect immediately without requiring an Association Update.

Verification

To verify the exchanged compression capability for a specific peer, use the following command:

```
show sx peers full address <peer-ip>
```

Sample Output:

```
Peer IP      : 21.21.21.101
Sx Service Id: 1
Group Name   : default
...
Compression  : Enabled
```

Configuring N4 Interface

This section describes the following configurations:

- Identifying N4 Interface
- Adding N4-type and Modification of N4 Parameters in Sx Service

Identifying an N4 Interface

Use the following configuration to identify if a peer node is an N4 interface type.

```
configure
   control-plane-group group_name
      peer-node-id [ ipv4-address ipv4_address | ipv6-address ipv6_address ]
interface n4
      end
```

NOTES:

- To enable the **n4 interface** CLI command, you need the **require upf** CLI command on the UPF, which depends on the UPF license.
- **[ipv4-address** *ipv4_address* **| ipv6-address** *ipv6_address* **]** :

  ipv4-address *ipv4_address*: Specifies the IPv4 address of the peer node.

  ipv6-address *ipv6_address*: Specifies the IPv6 address of the peer node.
- **interface n4**: Identifies the N4 interface.

Modification of N4-type Parameters in an Sx Service

Use the following configuration to modify N4-type parameters in an Sx Service.

```
configure
   context context_name
      sx-service service_name
         n4 [ max-retransmissions max_retransmission_value |
retransmission-timeout-ms timeout_value ]
         end
```

NOTES:

- **n4**: Allows modifications to N4 parameters.
- **[max-retransmissions***max_retransmission_value* **| retransmission-timeout-ms** *timeout_value***]**:

  max-retransmissions *max_retransmission_value* Configures maximum retries for Sx control packets. *max_retransmission_value* must be an integer in the range of 0 to 15. The default value is 4.

  retransmission-timeout-ms: Configures the control packet retransmission timeout in Sx in milliseconds. *timeout_value* must be an integer in the range of 1000 to 20000 milliseconds. The timeout value must be configured in steps of 100; for example: 1000, 1100, 1200, and so on. The default value is 5000 milliseconds.

Statistics

This section provides information on show commands and their output available in support of this feature.

show control-plane-group

The output of this command displays the following fields for this feature:

- Interface Type – This field indicates if the peer interface is N4. It is not displayed for non-N4 interfaces.

show sx-service all

The output of this command displays the following fields for this feature:

- N4
  - N4 Retransmission Timeout
  - N4 Maximum Request Retransmission

show subscribers user-plane-only all

The output of this command displays the following fields for this feature:

- Interface
  - N4

show user-plane-service statistics all

The output of this command displays the following fields for this feature:

- N4 interface-type PDNs
  - Active
  - Setup
  - Released

show subscribers user-plane-only seid number pdr all

The output of this command displays the following fields for this feature:

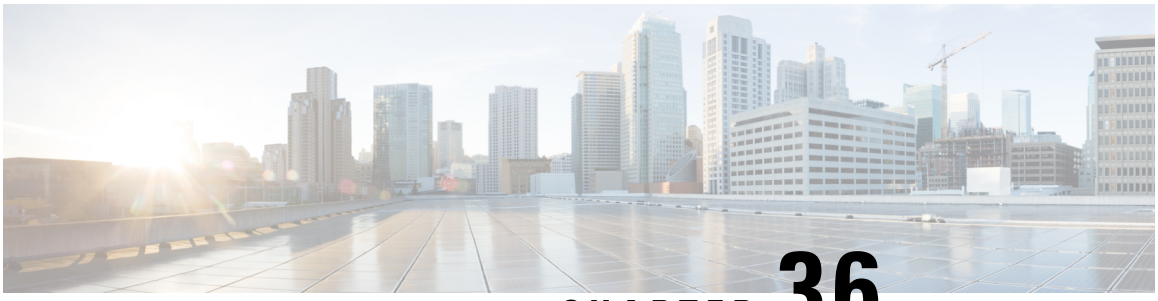
- Associated-QFIs

show subscribers user-plane-only callid number pdr full all

The output of this command displays the following fields for this feature:

- QoS Flow Identifier

show subscribers user-plane-only callid number pdr full all



CHAPTER 36

N4/Sx over IPSec

- Feature Summary and Revision History, on page 367
- Feature Description, on page 368
- Recommended Timers, on page 370
- Sample Configurations, on page 377
- Monitoring and Troubleshooting, on page 381

Feature Summary and Revision History

Summary Data

Table 81: Summary Data

Applicable Product(s) or Functional Area	5G-UPF
Applicable Platform(s)	VPC-SI SMI
Feature Default Setting	Enabled – Always-on
Related Changes in this Release	Not Applicable
Related Documentation	*UCC 5G UPF Configuration and Administration Guide*

Revision History

Revision Details	Release
First introduced.	2022.04.0

Feature Description

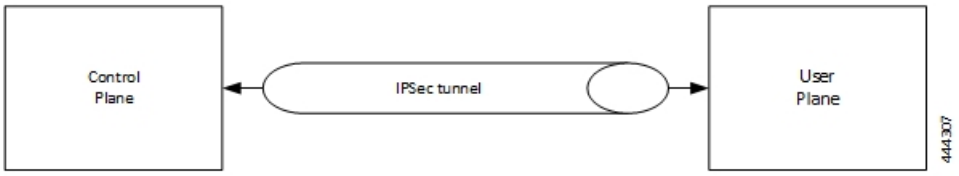
Internet Protocol Security (IPSec) is a suite of protocols that interact with one another to provide secure private communication across IP networks. These protocols allow the system to establish and maintain secure tunnels with peer security gateways. IPSec provides confidentiality, data integrity, access control, and data source authentication to IP datagrams.

In Cisco Cloud Native 5G, the IPSec functionality is available in Tunnel mode both on Session Management Function (SMF) and User Plane Function (UPF). The IPSec crypto-maps are associated under the appropriate interface on respective nodes. The IPSec tunnel is created between each SMF or UPF pair explicitly. This feature supports the IPv4 and IPv6 tunneling mode. There is no change on the N4/Sx service configuration.

The IPSec tunnel mode encapsulates the entire IP packet to provide a virtual secure hop between two gateways. It forms VPN kind of functionality, where entire IP packets are encapsulated inside another and delivered to the destination. It encapsulates the full IP header as well as the payload.

The N4 over IPSec feature supports TLS 1.2 using OpenSSL certificates.

Figure 24: N4/Sx Over IPSec Tunnel



When N4/Sx over IPSec is enabled on UPF NF running VPP, then the following parameter must be used under "VPP Param" for the N4/Sx Over IPSec feature to work.

VPP_DPDK_DATA_SIZE=5120

The VPP Param is stored in the **staros_para.cfg** file on a CD-ROM and this configuration is read and applied to VPP by UPF during its boot.

Note This parameter is supported until VPP version 19.08. This parameter introduces a memory overhead of about 800 MB. You must consider this condition before using the feature. If the UPF has less RAM, then VM must be allocated with extra 1 GB of RAM memory for the feature to work properly.

For more information on IPSec support, see the StarOS *IPSec Reference*.

IKEv2 Keep-Alive Messages (Dead Peer Detection)

IPSec for N4/Sx interface supports IKEv2 keep-alive messages, also known as Dead Peer Detection (DPD), originating from both ends of an IPSec tunnel. As per RFC 3706, DPD is used to simplify the messaging required to verify communication between peers and tunnel availability.

IPSec DPD is an optional configuration. If its disabled, the IPSec node does not initiate DPD request. However, the node always responds to DPD availability messages initiated by peer node regardless of its DPD configuration.

The following method can be used to calculate the keep-alive interval value when N4/Sx over IPSec feature is configured:

```
((max-retransmissions + 1) * retransmission-timeout-ms) * 2
```

The keep-alive interval value specifies the time that the IPSec tunnel will remain up till DPD is triggered.

Example:

The following is a sample output of the **show configuration context** *context_name* **verbose** CLI command under N4/Sx service:

```
sx-service sx
   instance-type userplane
   bind ipv4-address 192.168.1.1 ipv6-address bbbb:abcd::11
   sxa max-retransmissions 4
   sxa retransmission-timeout-ms 5000
```

Here, the value of **max-retransmissions** is 4 and **retransmission-timeout-ms** is 5000. Therefore, the keep-alive interval value will be 50:

((max-retransmissions + 1) * retransmission-timeout-ms) * 2 = Keep-alive interval

((4+1) * 5000) * 2 = 50

IKESA Rekey

UPF supports both IKESA Rekey and IPSec Rekey.

For IKESA Rekey, the **lifetime** *interval* CLI must be configured under **ikev2-ikesa transform-set** *transform_set*. You must also configure **ikev2-ikesa rekey** under **crypto map** configuration. Following is a configuration example:

```
ikev2-ikesa transform-set ikesa-foo
   encryption aes-cbc-256
   group 14
   hmac sha2-256-128
   lifetime 28800
   prf sha2-256
...
...
...
crypto map foo0 ikev2-ipv4
   match address foo0
   authentication local pre-shared-key encrypted key secret_key
   authentication remote pre-shared-key encrypted key secret_key
   ikev2-ikesa max-retransmission 3
   ikev2-ikesa retransmission-timeout 15000
   ikev2-ikesa transform-set list ikesa-foo
   ikev2-ikesa rekey
   keepalive interval 50
   control-dont-fragment clear-bit
   payload foo-sa0 match ipv4
      ipsec transform-set list A-foo
      lifetime 600
      rekey keepalive
   #exit
   peer 172.19.222.2
   ikev2-ikesa policy error-notification
```

Limitations

The feature is supported only in IPv4-IPv4 tunneling mode.

Recommended Timers

The following table provides the recommended timer values for CLI commands related to IPSec, N4/Sx, and SRP.

IPSec	**SMF**	**UPF**
ikev2-ikesa max-retransmission	*3*	*3*
ikev2-ikesa retransmission-timeout	*1000*	*1000*
keepalive	**interval** *4* **timeout** *1* **num-retry** *4*	**interval** *5* **timeout** *2* **num-retry** *4*
N4/Sx	**SMF**	**UPF**
sx-protocol heartbeat interval	*10*	*10*
sx-protocol heartbeat retransmission-timeout	*5*	*5*
sx-protocol heartbeat max-retransmissions	*4*	*4*
sxa max-retransmissions	*4*	*4*
sxa retransmission-timeout-ms	*5000*	*5000*
sxb max-retransmissions	*4*	*4*
sxb retransmission-timeout-ms	*5000*	*5000*
sxab max-retransmissions	*4*	*4*
sxab retransmission-timeout-ms	*5000*	*5000*
sx-protocol association reattempt-timeout	*60*	*60*
SRP	**SMF**	**UPF**
hello-interval	*3*	*3*
dead-interval	*15*	*15*

Recommended Configurations

Following are the recommended configurations and restrictions related to N4/Sx and SRP over IPSec:

- The multihop BFD timer between SMF and UPF must be seven seconds (for Data UPFs).

- The singlehop BFD must be enabled on all the contexts (SMF GW/Billing and UPF Gn/Gi).
- Inter-chassis multihop BFD must be enabled for SMF-SMF ICSR and UPF-UPF ICSR (IMS UPF).
- The SRP-IPSec ACL must be configured for TCP protocol instead of IP protocol.
- The N4/Sx-IPSec ACL must be configured for UDP protocol instead of IP protocol.

Example Configurations in SMF

Multihop BFD Configuration VPC-DI

The following is an example of multihop BFD configuration with seven seconds timer.

```
bfd-protocol
      bfd multihop-peer 209.165.200.226 interval 350 min_rx 350 multiplier 20
      bfd multihop-peer 209.165.200.227 interval 350 min_rx 350 multiplier 20
      bfd multihop-peer 209.165.200.225 interval 350 min_rx 350 multiplier 20
      bfd multihop-peer 209.165.200.230 interval 350 min_rx 350 multiplier 20
      bfd multihop-peer 209.165.200.228 interval 350 min_rx 350 multiplier 20
      bfd multihop-peer 209.165.200.229 interval 350 min_rx 350 multiplier 20
    #exit
```

Multihop BFD Configuration VPC-SI

The following is an example of multihop BFD configuration with three seconds timer.

```
bfd-protocol
      bfd multihop-peer 209.165.200.226 interval 150 min_rx 150 multiplier 20
      bfd multihop-peer 209.165.200.227 interval 150 min_rx 150 multiplier 20
      bfd multihop-peer 209.165.200.225 interval 150 min_rx 150 multiplier 20
      bfd multihop-peer 209.165.200.230 interval 150 min_rx 150 multiplier 20
      bfd multihop-peer 209.165.200.228 interval 150 min_rx 150 multiplier 20
      bfd multihop-peer 209.165.200.229 interval 150 min_rx 150 multiplier 20
    #exit
```

BGP Configuration

The following is an example of BGP configuration with recommended timers.

```
router bgp 1111
      router-id 209.165.200.225
      maximum-paths ebgp 15
      neighbor 209.165.200.250 remote-as 1000
      neighbor 209.165.200.250 ebgp-multihop
      neighbor 209.165.200.250 update-source 209.165.200.225
      neighbor 1111:2222::101 remote-as 1000
      neighbor 1111:2222::101 ebgp-multihop
      neighbor 1111:2222::101 update-source 1111:2222::1
      bgp graceful-restart restart-time 120
      bgp graceful-restart stalepath-time 300
      timers bgp keepalive-interval 30 holdtime-interval 90 min-peer-holdtime-interval 0
server-sock-open-delay-period 10
      address-family ipv4
        redistribute connected
      #exit
      address-family ipv6
        neighbor 1111:2222::101 activate
        redistribute connected
      #exit
    #exit
```

Singlehop BFD Configuration

The following is an example of singlehop BFD configuration with three seconds timer.

```
interface bgp-sw1-2161-10
      ip address 209.165.200.233 209.165.200.255
      ipv6 address 1111:222::9/112 secondary
      bfd interval 999 min_rx 999 multiplier 3
    #exit
    interface bgp-sw1-2161-11
      ip address 209.165.200.234 209.165.200.255
      ipv6 address 1111:222::10/112 secondary
      bfd interval 999 min_rx 999 multiplier 3
    #exit
    interface bgp-sw1-2161-12
      ip address 209.165.200.235 209.165.200.255
      ipv6 address 1111:222::11/112 secondary
      bfd interval 999 min_rx 999 multiplier 3
    #exit
    interface bgp-sw1-2161-3
      ip address 209.165.200.226 209.165.200.255
      ipv6 address 1111:222::2/112 secondary
      bfd interval 999 min_rx 999 multiplier 3
    #exit
    interface bgp-sw1-2161-4
      ip address 209.165.200.227 209.165.200.255
      ipv6 address 1111:222::3/112 secondary
      bfd interval 999 min_rx 999 multiplier 3
    #exit
    interface bgp-sw1-2161-5
      ip address 209.165.200.228 209.165.200.255
      ipv6 address 1111:222::4/112 secondary
      bfd interval 999 min_rx 999 multiplier 3
    #exit
    interface bgp-sw1-2161-6
      ip address 209.165.200.229 209.165.200.255
      ipv6 address 1111:222::5/112 secondary
      bfd interval 999 min_rx 999 multiplier 3
    #exit
    interface bgp-sw1-2161-7
      ip address 209.165.200.230 209.165.200.255
      ipv6 address 1111:222::6/112 secondary
      bfd interval 999 min_rx 999 multiplier 3
    #exit
    interface bgp-sw1-2161-8
      ip address 209.165.200.231 209.165.200.255
      ipv6 address 1111:222::7/112 secondary
      bfd interval 999 min_rx 999 multiplier 3
    #exit
    interface bgp-sw1-2161-9
      ip address 209.165.200.232 209.165.200.255
      ipv6 address 1111:222::8/112 secondary
      bfd interval 999 min_rx 999 multiplier 3
    #exit
```

Static Route for Multihop BFD Configuration

The following is an example of static route multihop BFD configuration.

```
ip route static multihop bfd UP-5 209.165.200.240 209.165.200.245
    ip route static multihop bfd UP-6 209.165.200.240 209.165.200.246
    ip route static multihop bfd UP-9 209.165.200.240 209.165.200.247
    ip route static multihop bfd UP-10 209.165.200.240 209.165.200.248
    ip route static multihop bfd UP-7 209.165.200.240 209.165.200.249
    ip route static multihop bfd UP-8 209.165.200.240 209.165.200.250
```

Static Route for Singlehop BFD Configuration

The following is an example of static route singlehop BFD configuration.

```
ip route static bfd bgp-sw1-2161-3 209.165.200.230
    ip route static bfd bgp-sw1-2161-4 209.165.200.230
    ip route static bfd bgp-sw1-2161-5 209.165.200.230
    ip route static bfd bgp-sw1-2161-6 209.165.200.230
    ip route static bfd bgp-sw1-2161-7 209.165.200.230
    ip route static bfd bgp-sw1-2161-8 209.165.200.230
    ip route static bfd bgp-sw1-2161-9 209.165.200.230
    ip route static bfd bgp-sw1-2161-10 209.165.200.230
    ip route static bfd bgp-sw1-2161-11 209.165.200.230
    ip route static bfd bgp-sw1-2161-12 209.165.200.230
```

IPSec ACL Configuration

The following is an example IPSec ACL configuration in SMF.

```
ip access-list UP-1
      permit udp host 209.165.200.225 host 209.165.200.226
    #exit
```

IPSec Transform Set Configuration

The following is an example of IPSec Transform Set configuration in SMF.

```
ikev2-ikesa transform-set ikesa-UP-1
      encryption aes-cbc-256
      group 14
      hmac sha2-256-128
      lifetime 28800
      prf sha2-256

    ipsec transform-set A-UP-1
      encryption aes-cbc-256
      hmac sha2-256-128
      group 14
```

IPSec Crypto Map Configuration

The following is an example of IPSec Crypto Map configuration in SMF.

```
crypto map UP-1 ikev2-ipv4
      match address UP-1
      authentication local pre-shared-key encrypted key secretkey
      authentication remote pre-shared-key encrypted key secretkey
      ikev2-ikesa max-retransmission 3
      ikev2-ikesa retransmission-timeout 1000
      ikev2-ikesa transform-set list ikesa-UP-1
      ikev2-ikesa rekey
      keepalive interval 4 timeout 1 num-retry 4
      control-dont-fragment clear-bit
      payload foo-sa0 match ipv4
        ipsec transform-set list A-UP-1
        lifetime 300
        rekey keepalive
      #exit
      peer 209.165.200.224
      ikev2-ikesa policy error-notification
    #exit
```

N4/Sx Configuration

The following is an example of N4/Sx configuration in SMF.

```
sx-service SX-1
      instance-type controlplane
      sxa max-retransmissions 4
      sxa retransmission-timeout-ms 5000
      sxb max-retransmissions 4
      sxb retransmission-timeout-ms 5000
      sxab max-retransmissions 4
      sxab retransmission-timeout-ms 5000
      n4 max-retransmissions 4
      n4 retransmission-timeout-ms 5000
      sx-protocol heartbeat interval 10
      sx-protocol heartbeat retransmission-timeout 5
      sx-protocol heartbeat max-retransmissions 4
      sx-protocol compression
      sx-protocol supported-features load-control
      sx-protocol supported-features overload-control
    exit
end
```

Example Router Configurations

Static Routes for Interface

The following is an example configuration of static route for interface.

```
ip route 209.165.200.224/27 Vlan1111 209.165.200.225
ip route 209.165.200.224/27 Vlan1111 209.165.200.226
ip route 209.165.200.224/27 Vlan1111 209.165.200.227
ip route 209.165.200.224/27 Vlan1111 209.165.200.228
ip route 209.165.200.224/27 Vlan1111 209.165.200.229
ip route 209.165.200.224/27 Vlan1111 209.165.200.230
ip route 209.165.200.224/27 Vlan1111 209.165.200.231
ip route 209.165.200.224/27 Vlan1111 209.165.200.232
ip route 209.165.200.224/27 Vlan1111 209.165.200.233
ip route 209.165.200.224/27 Vlan1111 209.165.200.234
```

Static Routes for Singlehop BFD

The following is an example configuration of static route for singlehop BFD.

```
ip route static bfd Vlan1111 209.165.200.225
ip route static bfd Vlan1111 209.165.200.226
ip route static bfd Vlan1111 209.165.200.227
ip route static bfd Vlan1111 209.165.200.228
ip route static bfd Vlan1111 209.165.200.229
ip route static bfd Vlan1111 209.165.200.230
ip route static bfd Vlan1111 209.165.200.231
ip route static bfd Vlan1111 209.165.200.232
ip route static bfd Vlan1111 209.165.200.233
ip route static bfd Vlan1111 209.165.200.234
```

Interface for Singlehop BFD

The following is an example configuration of interface for singlehop BFD.

```
interface Vlan1111
  no shutdown
  bandwidth 10000000
  bfd interval 999 min_rx 999 multiplier 3
  no bfd echo
```

```
  ip address 209.165.200.224/27
  ipv6 address 1111:222::1/112
```

BGP Configuration

The following is an example of BGP configuration with recommended timers.

```
router bgp 1000
  router-id 209.165.200.226
  timers bgp 30 90
  timers bestpath-limit 300
  timers prefix-peer-timeout 30
  timers prefix-peer-wait 90
  graceful-restart
  graceful-restart restart-time 120
  graceful-restart stalepath-time 300
```

Example Configurations in UPF

IPSec ACL Configuration

The following is an example of IPSec ACL configuration in UPF.

```
ip access-list CP-1
      permit udp host 209.165.200.225 host 209.165.200.226
    #exit
```

IPSec Transform Set Configuration

The following is an example of IPSec Transform Set configuration in UPF.

```
ipsec transform-set A-CP-1
      encryption aes-cbc-256
      hmac sha2-256-128
      group 14

    ikev2-ikesa transform-set ikesa-CP-1
      encryption aes-cbc-256
      group 14
      hmac sha2-256-128
      lifetime 28800
      prf sha2-256
```

IPSec Crypto Map Configuration

The following is an example of IPSec Crypto Map configuration in UPF.

```
crypto map CP-1 ikev2-ipv4
      match address CP-1
      authentication local pre-shared-key encrypted key secretkey
      authentication remote pre-shared-key encrypted key secretkey
      ikev2-ikesa max-retransmission 3
      ikev2-ikesa retransmission-timeout 1000
      ikev2-ikesa transform-set list ikesa-CP-1
      ikev2-ikesa rekey
      keepalive interval 5 timeout 2 num-retry 4
      control-dont-fragment clear-bit
      payload foo-sa0 match ipv4
        ipsec transform-set list A-CP-1
      #exit
      peer 209.165.200.230
      ikev2-ikesa policy error-notification
    #exit
```

N4/Sx Configuration

The following is an example of N4/Sx configuration in UPF.

```
sx-service SX-1
      instance-type userplane
      sxa max-retransmissions 4
      sxa retransmission-timeout-ms 5000
      sxb max-retransmissions 4
      sxb retransmission-timeout-ms 5000
      sxab max-retransmissions 4
      sxab retransmission-timeout-ms 5000
      n4 max-retransmissions 4
      n4 retransmission-timeout-ms 5000
      sx-protocol heartbeat interval 10
      sx-protocol heartbeat retransmission-timeout 5
      sx-protocol heartbeat max-retransmissions 4
      sx-protocol compression
    exit
```

Example SRP Configurations

IPSec ACL Configuration

The following is an example of IPSec ACL configuration for SRP.

```
ip access-list SRP
      permit tcp host 209.165.200.227 host 209.165.200.228
    #exit
```

SRP Configuration

The following is an example of SRP configuration.

```
configure
  context srp
    bfd-protocol
      bfd multihop-peer 209.165.200.225 interval 999 min_rx 999 multiplier 3
    #exit
configure
  context srp
    service-redundancy-protocol
      chassis-mode primary
      hello-interval 3
      dead-interval 15
      monitor bfd context srp 209.165.200.226 chassis-to-chassis
      monitor bgp context gi-pgw 209.165.200.245
      monitor bgp context gi-pgw 3333:888::1
      monitor bgp context saegw 209.165.200.245
      monitor bgp context saegw 3333:888::2
      peer-ip-address 209.165.200.227
      bind address 209.165.200.228
    #exit
    ip route static multihop bfd srp 209.165.200.229 209.165.200.245
    ip route 209.165.201.1 209.165.202.129 209.165.200.230 SRP-Physical-2102
    ip route 209.165.201.2 209.165.202.130 209.165.200.231 SRP-Physical-2102
    ip route 209.165.201.3 209.165.202.131 209.165.200.232 SRP-Physical-2102
    ip igmp profile default
    #exit
  #exit
end
```

Sample Configurations

In following sample configuration, the N4/Sx and IPSec interface IP Addresses are defined as:

```
SMF N4/Sx - 192.0.2.1
UPF N4/Sx - 192.0.2.7
SMF IPSec - 198.51.100.1
UPF IPSec - 198.51.100.2
```



Note

- For this release, following are the recommended timer values on UPF:

  ```
  sx-protocol heartbeat retransmission-timeout 20
  sx-protocol heartbeat max-retransmissions 3
  ```

- For this release, following are the recommended timer values on SMF:

  ```
  sx-protocol heartbeat retransmission-timeout 20
  sx-protocol heartbeat max-retransmissions 5
  ```

Control Plane

IPSec Configuration

```
config
  context EPC-CP
    ip access-list foo0
      permit ip host 192.0.2.1 host 192.0.2.7
    #exit
    ipsec transform-set A-foo
    #exit
    ikev2-ikesa transform-set ikesa-foo
    #exit
    crypto map foo0 ikev2-ipv4
      match address foo0
      authentication local pre-shared-key key secret
      authentication remote pre-shared-key key secret
      ikev2-ikesa max-retransmission 3
      ikev2-ikesa retransmission-timeout 15000
      ikev2-ikesa notify-msg-error no-apn-subscription backoff-timer 0
      ikev2-ikesa notify-msg-error network-failure backoff-timer 0
      ikev2-ikesa transform-set list ikesa-foo
      ikev2-ikesa configuration-attribute p-cscf-v6 private length 0
      ikev2-ikesa configuration-attribute p-cscf-v6 iana length 0
      keepalive interval 50
      payload foo-sa0 match ipv4
        ipsec transform-set list A-foo
        lifetime 300
        rekey keepalive
      #exit
      peer 198.51.100.2
      ikev2-ikesa policy error-notification
      notify-payload error-message-type ue base 0
      notify-payload error-message-type network-transient-minor base 0
      notify-payload error-message-type network-transient-major base 0
      notify-payload error-message-type network-permanent base 0
    #exit
    interface CP_IPSEC loopback
      ip address 198.51.100.1 255.255.255.0
```

```
   crypto-map foo0
    #exit
end
```

N4/Sx Configuration

```
   sx-service SX-1
      instance-type controlplane
      bind ipv4-address 192.0.2.1
      sx-protocol heartbeat retransmission-timeout 20
      sx-protocol heartbeat max-retransmissions 5
   exit
```

User Plane

IPSec Configuration

```
config
  context EPC-UP
    ip access-list foo0
      permit ip host 192.0.2.7 host 192.0.2.1
    #exit
    ipsec transform-set A-foo
    #exit
    ikev2-ikesa transform-set ikesa-foo
    #exit
    crypto map foo0 ikev2-ipv4
      match address foo0
      authentication local pre-shared-key key secret
      authentication remote pre-shared-key key secret
      ikev2-ikesa max-retransmission 3
      ikev2-ikesa retransmission-timeout 15000
      ikev2-ikesa notify-msg-error no-apn-subscription backoff-timer 0
      ikev2-ikesa notify-msg-error network-failure backoff-timer 0
      ikev2-ikesa transform-set list ikesa-foo
      ikev2-ikesa configuration-attribute p-cscf-v6 private length 0
      ikev2-ikesa configuration-attribute p-cscf-v6 iana length 0
      keepalive interval 50
      payload foo-sa0 match ipv4
        ipsec transform-set list A-foo
      #exit
      peer 198.51.100.1
      ikev2-ikesa policy error-notification
      notify-payload error-message-type ue base 0
      notify-payload error-message-type network-transient-minor base 0
      notify-payload error-message-type network-transient-major base 0
      notify-payload error-message-type network-permanent base 0
    #exit
    interface UP_IPSEC  loopback
      ip address 198.51.100.2 255.255.255.0
   crypto-map foo0
    #exit
end
```

N4/Sx Configuration

```
   sx-service SX-1
      instance-type userplane
      bind ipv4-address 192.0.2.7 ipv6-address dddd:51:31:1:209::
      sxa max-retransmissions 12
      sxb max-retransmissions 12
      sxab max-retransmissions 12
      sx-protocol heartbeat interval 30
      sx-protocol heartbeat retransmission-timeout 20
```

```
      sx-protocol heartbeat max-retransmissions 3
    exit
```

SMI strongSwan Configuration

The following is a sample configuration to spawn the SMI strongSwan pod:

```
SMI:
addons strongswan enabled
strongswan connections N4_IPSec_RCM3
  auto add
  keyexchange ikev2
  type tunnel
  left 192.12.31.202
  right 50.50.29.5
  leftsubnet 192.12.31.202/24
  rightsubnet 50.50.29.4/32
  leftauth psk
  rightauth psk
  leftsendcert never
  psk starent
  esp aes128-sha1,aes128-sha256-prfsha256
  ike aes128-sha1-modp1024,aes128-sha256-modp1536
  reauth no
  dpdaction clear
  dpddelay 300
  dpdtimeout 60
  closeaction none
  server-cert "-----BEGIN CERTIFICATE-----MIIDjTCCAnWgAwIBAgIUF6njegbcarj2oq
/x9c2+utqPThUwDQYJKoZIhvcNAQELBQAwRTELMAkGA1UEBhMCQVUxEzARBgNVBAgMClNvbWUtU3R
hdGUxITAfBgNVBAoMGEludGVybmV0IFdpZGdpdHMgUHR5IEx0ZDAeFw0yMjA5MDcwOTQ0MDdaFw0zM
jA5MDQwOTQ0MDdaMEUxCzAJBgNVBAYTAkFVMRMwEQYDVQQIDApTb211LVN0YXRlMSEwHwYDVQQKDBh
JbnRlcm5ldCBXaWRnaXRzIFB0eSBMdGQwggEiMA0GCSqGSIb3DQEBAQUAA4IBDwAwggEKAoIBAQDkKA
vGj94OWcFV8j7Enpr5HHqQxakb7hD0fETPByMIb9lPA73AM/3g7YjyIuAFhhs/fx4ZbFQJKDUVjiK/
PE7Mq/Opw5vIsUAgyhors2goa3YvBEPCmTk4fPz21hkWLHZgTARKq3XkgdCAO7kB7UsJpxVBSGg0A
52bIy3bB5C8YNa4rTrafVqzzFdYrQfAama2lpLrfxI7TzoZ6qK1LUDe8U7K/Ln/LJOeqxXClGSEzz
GRBqG41FeU18u3mpJ1pDINUJj7E7r+UN58aTwMoW3/ThCL/2ou+vjTVN7TDzva6XdJPNBCMA5dKEh
0EF10rMo8nmtLzo4UW9NBKMbiv7KPAgMBAAGjdTBzMB8GA1UdIwQYMBaAFC/Lvz0LAowgIkydSpKNUwy/
wHzJMAkGA1UdEwQCMAAwCwYDVR0PBAQDAgTwMDgGA1UdEQQxMC+CFWNuZHAtbmFybWFkYS1tYXN0ZXIt
M4cEwAwfyocQIAFIiAGSEjEAAAAAAAACAjANBgkqhkiG9w0BAQsFAAOCAQEAB6WUQI4qgEHQ8E5sYwzP
zw5KC/zGP2WZIkBfcs8ReiGmLJlC8n8uceWH12ZbFwY75j3EBFfqkmnNXftQGmuU8oGyZsuPDpmEySo+
nE28xnQDZDGzABLZWLSZqqeR6obnYUKvDho14kd40o1hnVlaONw1mrwc/QyFvn3tOwoYnXgaktGM01Fu
cQY1Kc33DvJx3n7fOsdoOLRm9jEENYT3Dv8b6/Ezr2mMHRhAwuoaFpvOSc/eLJy0QO7RpQLpHcRmnh9n
XO+gccB+e0YzvuBS5PONt8wjNSCKl46ZW4F9jpvehR8P/rvH/3VbwDaa8c6xARHNxzNcfq5S4tK/f58RSA
==-----END CERTIFICATE-----"
  server-priv-key "$8$gFVXFkFlJplgshiCqWs222+/vkNL5suwjGgqQVwhm1HEIvNp5ViKE8Stz7NK
jubZL1uXIDuy\nTbZmSPp8gIyWFTAJadMNjSoJswWhFYHX+aYoliCIdWQEUFSnJTz2Gofjgex3kM7g8iFkw
BNb\nB6qnSOV2WwMHown1ZfIGEZQAZ6B6iNnQbIHVrOgsyAY6akkyoNzIuc1gFdijQ2W56wW6tQR1\n5EpV
5zweW/Nr0RoOma+ZjpKY8L2VDW30SZ+VwbeTWexrVVfTbYifYYURekyIr6SbK4wFwt+3\nLhBUIr/6zvOlV
QBh4GVEheB5IJid0vHSI3N91sxX+VRaBodSKyw22HpC5BgWanarhkd1KfCT\nmoLzQ7+Nw0X0UfbKTLM5G8
IhXGxqccjl8Jb8nZf490MGx+XrYMkNcFNJ7ua7bxNhl1goTyUs\n4Wbw9hcviv9ZD4leTwnSlqnv5Yfr0ED
GJVrkW2zFv808fKdkJ2rO9T843u9DOrKrFo6XMPT2\n3JU9RL6zlI6bUMTHRqy2xLDfTtDBrD5jg3joJdD7
nkQfCW6cS79cXTBSLTc79p8otX8Jy56n\nkg1uDmQvdY+PgmbByvjQLrPkFr1BQ0C/g5F1uTPSiy2bNGr9l
QF8LfV8kakQMsi+FT0BJbil\nXHxw9pMu2p9srsZRmiBUw4PMq5nQ69jqDJweoNwzjqcJKBvIV1mvKIzHll
Ha0jOqA/FqASn0\nXzmKuZwG49c8qJaE5JBTLTjxeD7tG6A5XuewKynAYWnynT/0xP0mMDMcwEPdOt4e/L4
WJUOJ\nUn4EVo9EMOeG/eRzqILwAbeo2faQtY3HR7c5qMGgnBk903zIVsxl7SP4ujR0HuRw3zq7co5y\n6O
GSmu5Q79EeHezgxn+uiCsPSBwD3gkjvCerdBi8lKPptp//J+XyFA6MdgTbjzb+MxsDXszt\nyXaojBhn9t1
RxwVepAyVesm511JdH/IeDGIYY21Q2DT/k3RT490yKQSu2U2J3n49PCsEtHTQ\nbJo0WmoBVzkysE5kkL2R
MMD6PN+oV8eSqXJHkc1lAFhTpB+TqXcUI+QM0DsLdC1KOr7I5a6P\nl7jdyFFKlbPW8bOe2BB+bKA+5lOQ0
ygb9hlM76WdKmr8hhaimIuH6covaqISrFJJ0IvXcaWS\nhiatKxAq/KhkdczeM0WS6Z8PFlUwRoqgL8X8tn
v1C6tvJbUNLOPgTbHYjkflOyEeFHgXVlXi\nnzE4iAADsRTaMT/3G5YuPjk7+0lmtiZRKHXUPy7LyRwJHNZ
vkaEY+LALhA0ukMpH4DcdDibh\n16kPVUPvWZNZ2Mw3kILH5raqICdGYDDuW1SwCLBeV2pqMuaTzFiSPpLt
5AFXKtF5u9vA+VIC\nJcWP77XVbPTkbsnSBtxFy32RlZY5rx6hLEf/XsMnPAOJvprZvWuc7F+KzrexMmYAY
bJbKE3S\n8POaPet9r8+mPkQf+F5NQD7r3iz0iZ7Hj4IVzm5cvlo9yfatvm03cDplBhVAvsa5dTRuJCq+\
```

```
n0UMpf6PbcZI3vhVjIGm7iR+SVSVrq27+lGW76MpnpGwffm/gnyVvg97wl21LmPuool6vKljs\nc9DBybr
dOIwf6gkHkfwDPITGZEbc0SiH3AnIc8Z6HPiCqm1jLJ+2PfC6xnJdLRgkB8sJA1UC\n1VikR2YOvSR26Z6
PI5x7Nhq73jlRMr2N7cvrbBgfjmQyluHa0H/fnOYh4/D6Va8ROWCM4Ca3\nGOPeGn/oJAY1qogLSad33OI
DLsExvyh52x8KvrhdBCRRY5EabXa97XG0TtRTgt6NDd9NwZZ0\n2xxE06VUMS307UBAmyQn7vuezVqcHtv
3H9NnDFPRVLsnyreNo04VjZN6PHtqqOei/sLfL1GK\nVzvleeNGfSAvh1kmFh13f9p1jXgnTwt4iFErpaR
1lvk5K1RoF/+Sjo0HYhETvJFxA/yd/2I0\nZQe3ob/W4hBqI069yQjHbk+9L6kGWzQl3Tl8Lw1/YU/2AXS
zW8V9wCV0OhNLwQezt7a8EBm4\nX3CsPNVhhixhdvC/rSrXFPJnXy0mrcuCXhqLitWRA5VO6883Yry7ldP
uHzcVTyLCxYm0PWaT\nif4TQ6BxvT/gz7Ic6F3dO/QwMKoyeA7RoRR3XpnFcN1QMNTrF17jg17hDFjJwBO
dsq1gfau5\nUDeG6HihvcggYwnbkTprwaHflK/tsTCReNi+j/+ei+4DIe0f2vFgqaGjHdaa6qGkpSXks4L
R\nhyVd9/y+"
  nodes master-3
  exit
 exit
strongswan connections N4_IPSec_V6
  auto add
  keyexchange ikev2
  type tunnel
  left 2001:4888:192:1231::202
  right 2001:4888:50:50::22
  leftsubnet 2001:4888:192:1231::202/64
  rightsubnet 2001:4888:50:50::21/128
  leftauth psk
  rightauth psk
  leftsendcert never
  psk starent
  esp aes128-sha1,aes128-sha256-prfsha256
  ike aes128-sha1-modp1024,aes128-sha256-modp1536
  reauth no
  dpdaction clear
  dpddelay 300
  dpdtimeout 60
  closeaction none
  server-cert "-----BEGIN CERTIFICATE-----MIIDjTCCAnWgAwIBAgIUF6njegbcarj2oq
/x9c2+utqPThUwDQYJKoZIhvcNAQELBQAwRTELMAkGA1UEBhMCQVUxEzARBgNVBAgMClNvbWUtU3R
hdGUxITAfBgNVBAoMGEludGVybmV0IFdpZGdpdHMgUHR5IEx0ZDAeFw0yMjA5MDcwOTQ0MDdaFw0zM
jA5MDQwOTQ0MDdaMEUxCzAJBgNVBAYTAkFVMRMwEQYDVQQIDApTb21lLVN0YXRlMSEwHwYDVQQKDBh
JbnRlcm5ldCBXaWRnaXRzIFB0eSBMdGQwggEiMA0GCSqGSIb3DQEBAQUAA4IBDwAwggEKAoIBAQDkKA
vGj94OWcFV8j7Enpr5HHqQxakb7hD0fETPByMIb9lPA73AM/3g7YjyIuAFhhs/fx4ZbFQJKDUVjiK/
PE7Mq/Opw5vIsUAgyhors2goa3YvBEPCmTk4fPz21hkWLHZgTARKq3XkgdCAO7kB7UsJpxVBSGg0A
52bIy3bB5C8YNa4rTrafVqzzFdYrQfAama2lpLrfxI7TzoZ6qK1LUDe8U7K/Ln/LJOeqxXClGSEzz
GRBqG41FeU18u3mpJ1pDINUJj7E7r+UN58aTwMoW3/ThCL/2ou+vjTVN7TDzva6XdJPNBCMA5dKEh
0EF10rMo8nmtLzo4UW9NBKMbiv7KPAgMBAAGjdTBzMB8GA1UdIwQYMBaAFC/Lvz0LAowgIkydSpKNUwy/
wHzJMAkGA1UdEwQCMAAwCwYDVR0PBAQDAgTwMDgGA1UdEQQxMC+CFWNuZHAtbmFybWFkYS1tYXN0ZXIt
M4cEwAwfyocQIAFIiAGSEjEAAAAAAAACAjANBgkqhkiG9w0BAQsFAAOCAQEAB6WUQI4qgEHQ8E5sYwzP
zw5KC/zGP2WZIkBfcs8ReiGmLJlC8n8uceWH12ZbFwY75j3EBFfqkmnNXftQGmuU8oGyZsuPDpmEySo+
nE28xnQDZDGzABLZWLSZqqeR6obnYUKvDho14kd40o1hnVlaONw1mrwc/QyFvn3tOwoYnXgaktGM01Fu
cQY1Kc33DvJx3n7fOsdoOLRm9jEENYT3Dv8b6/Ezr2mMHRhAwuoaFpvOSc/eLJy0QO7RpQLpHcRmnh9n
XO+gccB+e0YzvuBS5PONt8wjNSCKl46ZW4F9jpvehR8P/rvH/3VbwDaa8c6xARHNxzNcfq5S4tK/f58RSA
==-----END CERTIFICATE-----"
  server-priv-key "$8$jW1TFB0/rJW4V6NrVjk4+1KE7Dw6ynkP3Bqtiwp2k+GQDrI4bX2n+a6Yvyeq
zzKdQ+EQLuy6\ncj9xOrxtNflmzaptNF9Ku786m934ID9hzmC8ISya6/4f2Xu+WdG6uIJl2jDhB/3B2PIc
b6VQ\nV7c4GmwPRNBlIZTVMvTS/2xiUx9bdXIQTVzl2Sc3bZqwLJ6ho/qr4r++T7blVZ16j5sYxUI6\nat
ZKNMMk8+0aoH4UaOd5vtoSkhXCLXkfyrGYagx4KceKxPxSciSEptAzM36py7hDqazW5epU\nFaAnw3PMhq
Ut1r790CaG3VZR5WpcJVkHbdpf0iMCt6pJjNeNlL7BTvns+vo16Mcgt0pyi6Rj\nBAo5lSzog9max0EiRk
spb4a91DFX8mV4tzTy0RCzbgkuzdZ3ecbB9OOvrkWOv7dLiWsZe66Q\nrQA4SLH7eOkgRQvDzqmx3DpqXP
7rebptkLAGXAxZV5uvUuyivdal0EPiB6fu3OwP+gJweZIg\nLjIVbJsREgN7YdukihOmk/xbSMK25Eu3X3
yI1Y55vvQfsY08WEfKBO+Alzjrvz4ABydVJcEE\nqywkSUK/j0VksGvN4lzgely07tpz22VjMTrxJvoWB+
5j9j183T/C1Wgf53miFz2z8ak0NYYe\n2AEP9NNs4tFkB9bY9JQsFv6zY3J+2hQ8iyCiYIROd5ItRyLenO
Bt1fKGp5FHg7dlPuOz0VoI\nUm7GlEexMIycNEr9rzOqzBbMiH5c53htY4iQWFvOARHhw2f5GWPZOIe8Z8
uTq4k5iUjWmaLa\nfhNMXIGX2QNgoduZwXiX6yv3gCpK8WDgF4dlvPjFB+f+iA+QyBlc5AYZuE/2yjYRCr
aaPkzx\ndrWm+Uh18d1YdmQq4ss/rUY4Q0DxDblv94Xx64NIq8dbnY7Zehjs9LXXHk2X6daSTC/FYIY+\n
J/Sks+tmnZ489Tojo/F5dS9iVVstP68MdKOl4OC53lDkqcNl0xhniu2nneS6HTlzUrKFSi4I\n3eRW0FwK
ONKrePxKdObZFB+FvV+xOa2UKnXBbpIh/ENFE0XnADP7Ljox3YsZZpvnXTcz7OcE\nq1b8ggtLt6KyWDb0
tdZljLAb4posj1NioJQzyxHT0gsdfkZxtWiqgt65gqXS/iDo9XXdFEj1\n1Gr07SoE+HwoJIPzAG/8fNhm
```

```
27CCYaUW1uWJAOUt9I5UCbER+2kFaVC+odEY2W5/Hfc/gWAy\nSRd6/46kvjN/SzabIdb1yqr68G4LIrHg
1kEBoAf9hXSkD7WY2SMSjl950Co4Cq6zVbk6PAcX\n2ET02FhFewiq+TamVNr0/ruyPohyr1Cpgjqzvx6s
+s7EMWOPJh+XEh8PPBKY+DcDEr2RBIcZ\n4O9uAzwiYTm1x4u8dw5kKRd+H8HFobgaQ18i3IfCdZ4DXyrq
mMrxOw52fGIuAd3Ln2j/AitZ\nQP61dlQlwWPHX7ykvXqCP6oZnfbMUVW3iIdFauZLiCDZkX98UXY3IZUi
Eq13GL6KZtKwFAbu\nqHYEtb5OJRiHRZnqmoOf7BQjC3cdDTBpmPd/s9JCSejqSahlSaSl4Qghbba35HIG
o9PVyks3\nGny6P6twYnDCHLdXbmfE+HE71MnRRPd63CNp+SeX/pP5nBhu6RU/K61ovPrSqcsmo8GiytZB
\ndtVH7nYk3ZUWan04us/bd5zNZfrQ1mCF4rS4KGprQ0N7xWSCilMI49aIxSCNM9WkYl8HAEWA\n1IceEZ
RxHxeltl3JMNTxwFlkP4i4IEalf//Tidrd13jka0NnnjOsboUgn7lay3LvsC6zsIGN\nkP0sTGI1hHj9OL
1iKkw6IhTS1Py5Bjof+XPE844QwOY6Qj6GTd5F/GQJOOD3rL2J/S651TvJ\nsnKM5roovBVb1dUANC/Eay
frpC/2w8wjqPQ/O02SzVNSbZOIPn8P8BV3Ql+NC8rEWL1FZMkI\nkM1AsZTx8BQ0Z1Haf4uhtV5+/29ula
EqEiTH1x2QDV9idWPekqr7eC3009YoGESWHuIH/JE/\n76Rr2zi4wK0JVecxbCGDOynIRFE3I3gRdxgtTi
GrOMe2WdqsUDvDkcijCVpHo0JS3jFVONR8\nxbEo7RpxdrQJ5Zr/u/vx1jnQV/bXTzwlkqoy1g9C3V1m4m
NrqYz62taNTF7+NEVyWC/cp5CU\n5SDuAs3JmFyLaRvyU5SsmDbzlyj+z3DUaByHWlWtC5+klwXYoZOKIy
8zNj+1Kzxosk1wiVX1\nqT4CoAKX"
```

To validate the IPSec tunnel CLI on the SMF protocol pod and validate the *ipsec.yaml* file on SMF, see the *Interfaces Support > N4 Interface* chapter for sample SMI strongSwan configuration.

For the latest strongSwan configurations, see the *Ultra Cloud Core Subscriber Microservices Infrastructure Operations Guide*.

Monitoring and Troubleshooting

This section contains the sample CLI command output of show commands for the N4/Sx over IPSec feature in both SMF and UPF.

show crypto ikev2-ikesa security-associations summary

```
I - Initiator
  R - Responder
Mgr                                                                    Lifetime
ID  VPN Local IPSec GW:Port   Remote IPSec GW:Port  State             /Remaining
=== === ==================== ==================== ================ ===========
54  2   192.168.170.55 :500   192.168.196.55 :500   AUTH_COMPLETE(I) 86400/16448

1 IKEv2 Security Association found in this context.
```

show crypto ipsec security-associations summary

```
+------ SA state:           (E) - Established
      |                         (P) - Partially Established
      |                         (N) - No SAs
      |
      |+----- Rekey/Keepalive:    (D) - Rekey Disabled
      ||                          (E) - Rekey Enabled/No Keepalive
      ||                          (K) - Rekey Enabled/Keepalive
      ||
      ||+---- Crypto Type:        (D) - Dynamic Map
      |||                         (I) - IKEv1 Map
      |||                         (J) - IKEv2 Map
      |||                         (M) - Manual Map
      |||                         (C) - CSCF Map
      |||
      |||
      VVV           Map Name                                                 Rekeys En Pkts
    De Pkts
===== === ============================================================== ====== ==========
 ==========
1     EDJ foo0                                                              0      3496
     3496
```

```
1 Crypto Map Found.
1 Crypto Map Established.
```

To validate the IPSec tunnel CLI on the SMF protocol pod and validate the *ipsec.yaml* file on SMF, see the *Interfaces Support > N4 Interface* chapter for sample SMI strongSwan configuration.

For the latest strongSwan configurations, see the *Ultra Cloud Core Subscriber Microservices Infrastructure Operations Guide*.

SMI strongSwan Validation

The following is a sample SMI strongSwan configuration to validate the IPSec tunnel CLI on SMF protocol pod:

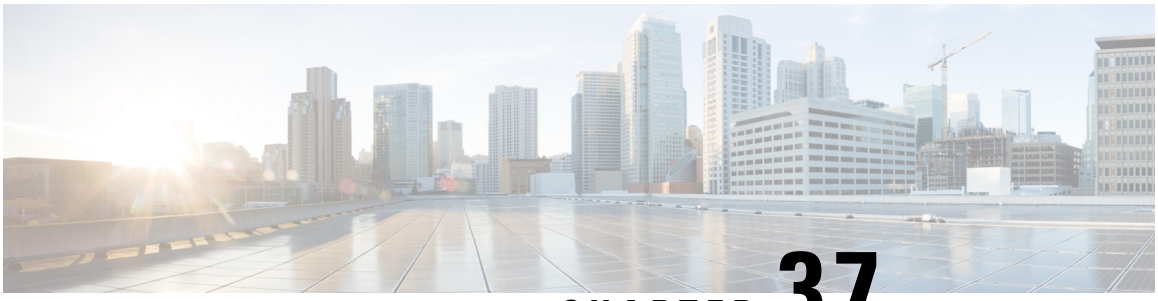
```
cloud-user@cndp-narmada-master-1:~$ kubectl exec -ti strongswan-d7zzp -n smi-strongswan --
 ipsec statusall
Status of IKE charon daemon (strongSwan 5.9.3, Linux 5.4.0-122-generic, x86_64):
uptime: 14 days, since Sep 19 16:36:02 2022
malloc: sbrk 6221824, mmap 0, used 3867536, free 2354288
worker threads: 11 of 16 idle, 5/0/0/0 working, job queue: 0/0/0/0, scheduled: 15
loaded plugins: charon aesni aes des rc2 sha2 sha1 md5 mgf1 random nonce x509 revocation
constraints pubkey pkcs1 pkcs7 pkcs8 pkcs12 pgp dnskey sshkey pem openssl af-alg fips-prf
gmp curve25519 xcbc cmac hmac ccm gcm drbg curl files attr kernel-netlink resolve
socket-default stroke vici updown eap-identity eap-dynamic eap-tls xauth-generic counters
Listening IP addresses:
10.105.90.166
2001:420:5504:2004::90:18
71.71.71.15
71.71.71.70
71.71.71.72
71.71.71.73
192.12.31.21
2001:4888:192:1231:42a6:b7ff:fe3b:7161
2001:4888:192:1231::21
192.12.31.203
192.12.31.206
192.12.31.207
192.12.31.208
192.12.31.209
192.12.31.210
192.12.31.211
192.12.31.213
192.12.31.214
192.12.31.215
2001:4888:192:1231::203
2001:4888:192:1231::206
2001:4888:192:1231::207
2001:4888:192:1231::208
2001:4888:192:1231::209
2001:4888:192:1231::210
2001:4888:192:1231::211
2001:4888:192:1231::213
2001:4888:192:1231::214
2001:4888:192:1231::215
2001:4888:192:1231::121
192.12.31.121
192.12.31.205
192.12.31.212
2001:4888:192:1231::205
2001:4888:192:1231::212
192.12.31.202
192.12.31.221
192.12.31.222
```

```
2001:4888:192:1231::202
2001:4888:192:1231::221
2001:4888:192:1231::222
192.50.0.1
fd00::1
192.115.3.37
Connections:
N4_IPSec_RCM1: 192.12.31.202...50.50.27.5 IKEv2, dpddelay=300s
N4_IPSec_RCM1: local: [192.12.31.202] uses pre-shared key authentication
N4_IPSec_RCM1: remote: [50.50.27.5] uses pre-shared key authentication
N4_IPSec_RCM1: child: 192.12.31.0/24 === 50.50.27.4/32 TUNNEL, dpdaction=clear
N4_IPSec_RCM2: 192.12.31.202...50.50.28.5 IKEv2, dpddelay=300s
N4_IPSec_RCM2: local: [192.12.31.202] uses pre-shared key authentication
N4_IPSec_RCM2: remote: [50.50.28.5] uses pre-shared key authentication
N4_IPSec_RCM2: child: 192.12.31.0/24 === 50.50.28.4/32 TUNNEL, dpdaction=clear
N4_IPSec_RCM3: 192.12.31.202...50.50.29.5 IKEv2, dpddelay=300s
N4_IPSec_RCM3: local: [192.12.31.202] uses pre-shared key authentication
N4_IPSec_RCM3: remote: [50.50.29.5] uses pre-shared key authentication
N4_IPSec_RCM3: child: 192.12.31.0/24 === 50.50.29.4/32 TUNNEL, dpdaction=clear
N4_IPSec_V6: 2001:4888:192:1231::202...2001:4888:50:50::22 IKEv2, dpddelay=300s
N4_IPSec_V6: local: [2001:4888:192:1231::202] uses pre-shared key authentication
N4_IPSec_V6: remote: [2001:4888:50:50::22] uses pre-shared key authentication
N4_IPSec_V6: child: 2001:4888:192:1231::/64 === 2001:4888:50:50::21/128 TUNNEL,
dpdaction=clear
N4_IPSec: 192.12.31.202...50.50.21.5 IKEv2, dpddelay=300s
N4_IPSec: local: [192.12.31.202] uses pre-shared key authentication
N4_IPSec: remote: [50.50.21.5] uses pre-shared key authentication
N4_IPSec: child: 192.12.31.0/24 === 50.50.21.4/32 TUNNEL, dpdaction=clear
Security Associations (5 up, 0 connecting):
N4_IPSec_RCM1[1345]: ESTABLISHED 96 minutes ago,
192.12.31.202[192.12.31.202]...50.50.27.5[50.50.27.5]
N4_IPSec_RCM1[1345]: IKEv2 SPIs: bc79a16793c7d7eb_i 087bb5cd20fd2f34_r*, rekeying in 72
minutes
N4_IPSec_RCM1[1345]: IKE proposal: AES_CBC_128/HMAC_SHA2_256_128/PRF_HMAC_SHA2_256/MODP_1536
N4_IPSec_RCM1{2501}: INSTALLED, TUNNEL, reqid 2, ESP SPIs: cd664851_i 13009213_o
N4_IPSec_RCM1{2501}: AES_CBC_128/HMAC_SHA2_256_128, 900 bytes_i (17 pkts, 16s ago), 829
bytes_o (17 pkts, 16s ago), rekeying in 36 minutes
N4_IPSec_RCM1{2501}: 192.12.31.202/32 === 50.50.27.4/32
N4_IPSec_RCM3[1343]: ESTABLISHED 97 minutes ago,
192.12.31.202[192.12.31.202]...50.50.29.5[50.50.29.5]
N4_IPSec_RCM3[1343]: IKEv2 SPIs: 1a50e1d11dfeacb7_i 7be50275473937a3_r*, rekeying in 65
minutes
N4_IPSec_RCM3[1343]: IKE proposal: AES_CBC_128/HMAC_SHA2_256_128/PRF_HMAC_SHA2_256/MODP_1536
N4_IPSec_RCM3{2499}: INSTALLED, TUNNEL, reqid 5, ESP SPIs: cd6f47ec_i 13009213_o
N4_IPSec_RCM3{2499}: AES_CBC_128/HMAC_SHA2_256_128, 1328 bytes_i (25 pkts, 26s ago), 1217
bytes_o (25 pkts, 26s ago), rekeying in 30 minutes
N4_IPSec_RCM3{2499}: 192.12.31.202/32 === 50.50.29.4/32
N4_IPSec_RCM2[1341]: ESTABLISHED 103 minutes ago,
192.12.31.202[192.12.31.202]...50.50.28.5[50.50.28.5]
N4_IPSec_RCM2[1341]: IKEv2 SPIs: 26fd8455c09927ab_i 78c5379f6559be4b_r*, rekeying in 60
minutes
N4_IPSec_RCM2[1341]: IKE proposal: AES_CBC_128/HMAC_SHA2_256_128/PRF_HMAC_SHA2_256/MODP_1536
N4_IPSec_RCM2{2500}: INSTALLED, TUNNEL, reqid 1, ESP SPIs: c6f5243c_i 13009213_o
N4_IPSec_RCM2{2500}: AES_CBC_128/HMAC_SHA2_256_128, 1026 bytes_i (19 pkts, 0s ago), 917
bytes_o (19 pkts, 0s ago), rekeying in 34 minutes
N4_IPSec_RCM2{2500}: 192.12.31.202/32 === 50.50.28.4/32
N4_IPSec_V6[1339]: ESTABLISHED 2 hours ago,
2001:4888:192:1231::202[2001:4888:192:1231::202]...2001:4888:50:50::22[2001:4888:50:50::22]
N4_IPSec_V6[1339]: IKEv2 SPIs: 64e5d5e102e885e7_i 9062914577d9eb95_r*, rekeying in 36 minutes
N4_IPSec_V6[1339]: IKE proposal: AES_CBC_128/HMAC_SHA2_256_128/PRF_HMAC_SHA2_256/MODP_1536
N4_IPSec_V6{2498}: INSTALLED, TUNNEL, reqid 3, ESP SPIs: c25a2531_i 0f009d13_o
N4_IPSec_V6{2498}: AES_CBC_128/HMAC_SHA2_256_128, 6817 bytes_i (89 pkts, 17s ago), 6326
bytes_o (89 pkts, 17s ago), rekeying in 17 minutes
N4_IPSec_V6{2498}: 2001:4888:192:1231::202/128 === 2001:4888:50:50::21/128
```

```
N4_IPSec[1337]: ESTABLISHED 2 hours ago, 192.12.31.202[192.12.31.202]...50.50.21.5[50.50.21.5]
N4_IPSec[1337]: IKEv2 SPIs: 9af4e1f24dcc0edb_i 6fcea88758803d37_r*, rekeying in 26 minutes
N4_IPSec[1337]: IKE proposal: AES_CBC_128/HMAC_SHA2_256_128/PRF_HMAC_SHA2_256/MODP_1536
N4_IPSec{2497}: INSTALLED, TUNNEL, reqid 4, ESP SPIs: cc7c3bf1_i 0f009c13_o
N4_IPSec{2497}: AES_CBC_128/HMAC_SHA2_256_128, 6844 bytes_i (121 pkts, 19s ago), 5693 bytes_o
 (121 pkts, 19s ago), rekeying in 4 minutes
N4_IPSec{2497}: 192.12.31.202/32 === 50.50.21.4/32
cloud-user@cndp-narmada-master-1:~$
```

The following is a sample SMI strongSwan configuration to validate the *ipsec.yaml* file on SMF:

```
cloud-user@cndp-narmada-master-1:~$ kubectl exec -ti strongswan-d7zzp -n smi-strongswan --
 cat /etc/ipsec.conf
conn N4_IPSec_RCM1
leftcert=/etc/ipsec.d/certs/N4_IPSec_RCM1.cert.pem
auto=add
closeaction=none
compress=no
dpdaction=clear
dpddelay=300
dpdtimeout=60
esp=aes128-sha1,aes128-sha256-prfsha256
ike=aes128-sha1-modp1024,aes128-sha256-modp1536
ikedscp=000000
ikelifetime=3h
keyexchange=ikev2
left=192.12.31.202
leftallowany=no
leftauth=psk
leftsendcert=never
leftsubnet=192.12.31.202/24
lifetime=1h
mobike=yes
reauth=no
rekey=yes
right=50.50.27.5
rightallowany=no
rightauth=psk
rightsubnet=50.50.27.4/32
sha256_96=no
type=tunnel
```



CHAPTER 37

N4 Session Management, Node Level, and Reporting Procedures

This chapter covers the following topics:

- Feature Summary and Revision History, on page 385
- Feature Description, on page 386
- How it Works, on page 387
- Configuring the N4 Session/Node Level Reporting Procedures, on page 397
- Enhanced PFCP Association Release Procedure for Graceful Session Termination, on page 401

Feature Summary and Revision History

Summary Data

Table 82: Summary Data

Applicable Product(s) or Functional Area	5G-UPF
Applicable Platform(s)	VPC-SI SMI
Feature Default Setting	Enabled – Always-on
Related Changes in this Release	Not Applicable
Related Documentation	Not Applicable

Revision History

Revision Details	Release
The **show user-plane-service statistics all** command has been enhanced to display PFCP IEs received in Sx Establishment or Modification messages over N4 interface.	2023.01 2021.02.2

Revision Details	Release
PDN Update procedure is introduced in this release.	2021.02.0
First introduced.	2020.02.0

Feature Description

N4 Session Management, Node Level, and Reporting Procedures

N4 Node-level Procedures

The N4 Node-level procedures in User Plane Function (UPF) involves the following processes:

- N4 Association Setup Procedure – The procedure used for setting up an N4 association between the Session Management Function (SMF) and UPF.
- N4 Association Update Procedure – The procedure used for modifying an existing N4 association between the SMF and UPF.
- N4 Association Release Procedure – The procedure used for terminating the N4 association between the SMF and UPF.
- N4 Heartbeat Procedure – The procedure used for sending and receiving the Heartbeat request and response.
- N4 Reporting Procedure – The procedure used for reporting echo request and response for the GTP-u path failure.

N4 Session Management

N4 session management procedures are used to control the functionality of the UPF. SMF can create, update, and remove the N4 session context in the UPF, which is described in 3GPP TS 23.501, clause 5.8.2.

The following procedures are performed in N4 Session Management:

- N4 Session Establishment
- N4 Session Modification
- N4 Session Deletion

NOTE: The SMF initiates all the above procedures.

N4 Session/Node-level Reporting Procedures

Whenever the data path between UPF and gNB is down, it is detected and reported to the SMF for corrective actions. The mechanism to detect and report it to SMF is clearly defined in 3GPP specifications. The reporting happens per GTP-u Tunnel level or per GTP-u endpoint level.

Relationships

The following features support the N4 session management, node level, and reporting procedures.

End Marker Support

The UPF sends the End Marker packets to support the reordering function in the target Radio Access Network (RAN). The UPF constructs the End Marker packets that are required for the reordering function.

Constructing the End Marker Packets through UPF

At the time of the handover procedure, the PDU session for the UE – which comprises of an UPF node – acts as a PDU session anchor and an intermediate UPF terminating N3 reference point. The SMF sends an N4 Session Modification Request message with the new AN Tunnel Info of NG RAN to specify the UPF to switch to the N3 paths. In addition, the SMF also specifies the UPF to send the End Marker packets on the old N3 user plane path.

After the UPF receives the indication, the End Markers are constructed and sent to each N3 GTP-U tunnel toward the source NG RAN, after sending the last PDU on the old path.

UEs IPv4, IPv6, and IPv4v6 Support

The UPF supports UE's IPv4, IPv6, and IPv4v6 sessions.

The N4 Session Establishment and Modification procedure for IPv6 sessions is the same as for IPv4 sessions. After the session is established, the SMF sends Router Advertisement (RA) message to UE announcing the IPv6 prefix to be used for traffic. Optionally, to get the IPv6 parameter from SMF faster, the UE can also initiate IPv6 Router Solicitation (RS).

The N4 Session Establishment and Modification procedure for IPv4v6 Session are similar to the IPv4 or IPv6 sessions except for the allocation of two UE IP addresses - one for IPv4 and the other for IPv6. The SMF sends the Router Advertisement message to the UE announcing the IPv6 prefix used for traffic after the session is established. Optionally, the UE can also initiate the IPv6 Router Solicitation to receive the IPv6 parameter from the SMF quickly.

How it Works

This section describes the N4 node-level, session management, and reporting procedures and associated call flows.

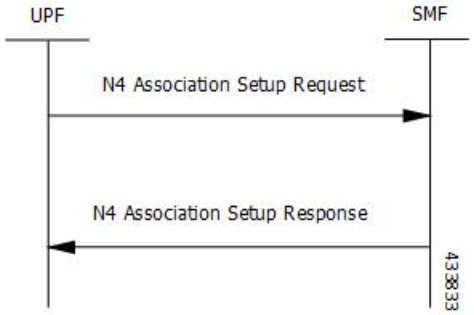
N4 Node-level Procedure Call Flows

N4 Association Setup Procedure Call Flow

The N4 Association Setup procedure creates the N4 association between the SMF and the UPF, which enables the SMF to use the UPF resources to establish N4 sessions. The N4 association setup procedure involves the following steps:

1. The UPF initiates the procedure by sending N4 Association Setup Request to the SMF.
2. The SMF sends an N4 Association Setup Response after it receives the request from the UPF.

The following call flow describes the UPF-initiated N4 Association Setup procedure:



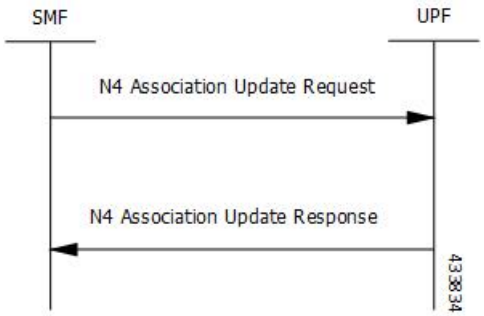
The UPF sends the following PFCP Association Setup Request:

- Node ID (UPF).
- Supported optional features in UPF. The UPF supports F-TEID allocation and release, sending of End Marker, and so on.

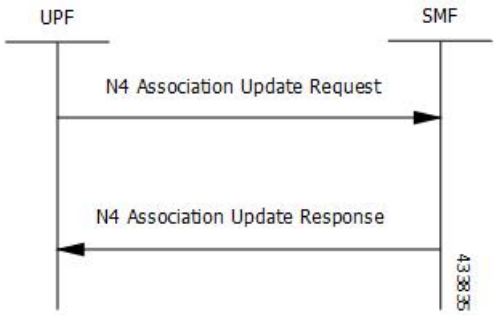
N4 Association Update Procedure Call Flow

The N4 Association Update procedure modifies an existing N4 association between the SMF and the UPF. It can be initiated either by the UPF or by the SMF to update the supported features or available UPF resources.

The following call flow depicts the SMF-initiated N4 Association Update procedure:



The following call flow depicts the UPF-initiated N4 Association Update procedure:



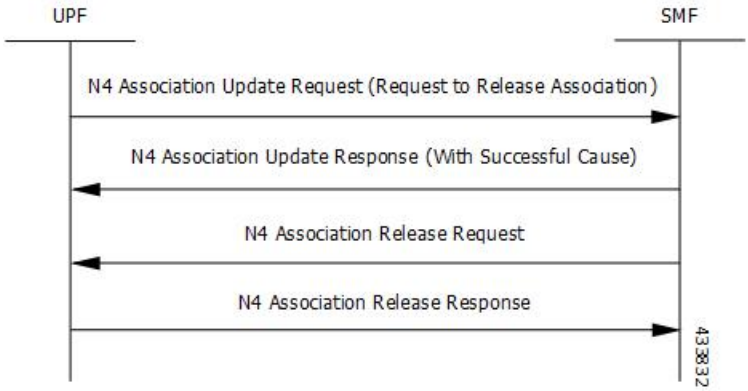
N4 Association Release Procedure Call Flow

The N4 Association Release procedure terminates the N4 association between the SMF and the UPF. It can be initiated either by the SMF or by the UPF. The UPF requests the SMF to perform the release of PFCP association by sending a PFCP Association Update Request. The SMF then initiates a PFCP Association Release Request to release the PFCP association.

The following call flow depicts the SMF-initiated N4 Association Release procedure:



The following call flow depicts the UPF-initiated N4 Association Release procedure:



N4 Heartbeat Procedure

The PFCP Heartbeat procedure includes the following messages:

- Heartbeat Request
- Heartbeat Response

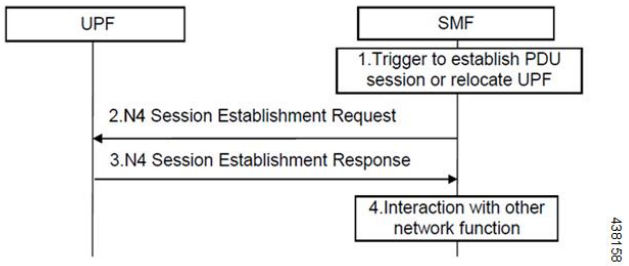
N4 Session Management Procedures Call Flows

The following section describes the N4 Session Management procedures.

N4 Session Establishment Call Flow

N4 Session Establishment is used to create the initial N4 session context for a PDU session at the UPF. SMF assigns a new N4 session ID and provides it to the UPF. The N4 session ID is stored by both entities and used

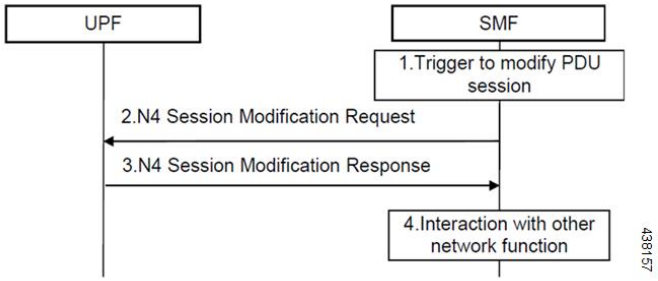
to identify the N4 session context during their interaction. SMF also stores the relation between the N4 session ID and PDU session for a UE.



Step	Description
1	SMF receives the trigger to establish a new PDU session or change the UPF for an established PDU session.
2	SMF sends an N4 session establishment request message to the UPF that contains the structured control information which defines how the UPF needs to behave.
3	UPF responds with an N4 session establishment response message containing any information that the UPF has to provide to the SMF in response to the control information received.
4	SMF interacts with the network function which triggered this procedure. For example, AMF or PCF.

N4 Session Modification Call Flow

N4 Session Modification is used to update the N4 session context of an existing PDU session at the UPF, which is executed between SMF and UPF whenever PDU session-related parameters have to be modified.

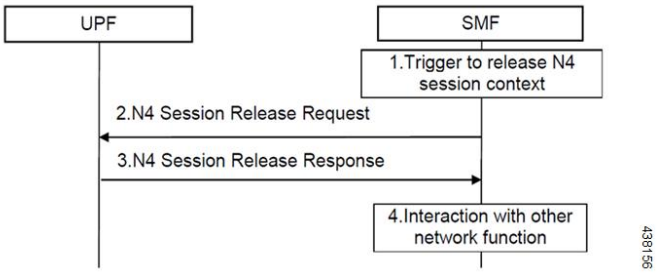


Step	Description
1	SMF receives the trigger to modify the existing PDU session.
2	SMF sends an N4 session modification request message to the UPF which contains the update for the structured control information that defines how the UPF needs to behave.

Step	Description
3	UPF identifies the N4 session context to be modified by the N4 session ID and updates the parameters of this N4 session context according to the list of parameters that are sent by the SMF. UPF responds with an N4 session modification response message containing any information that the UPF has to provide to the SMF in response to the control information received.
4	SMF interacts with the network entity which triggered this procedure. For example, AMF or PCF.

N4 Session Delete Call Flow

N4 Session Delete is used to remove the N4 session context of an existing PDU session at the UPF.

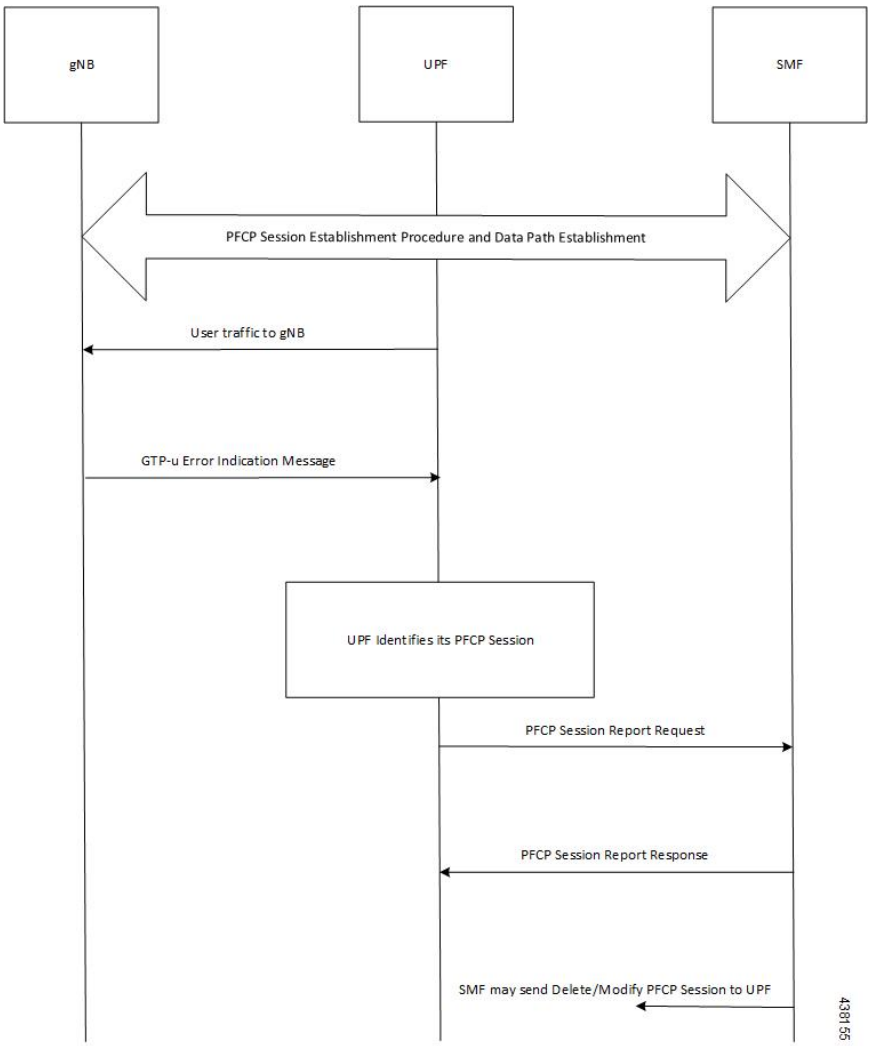


Step	Description
1	SMF receives the trigger to remove the N4 session context for the PDU session.
2	SMF sends an N4 session delete request message to the UPF.
3	UPF identifies the N4 session context to be removed by the N4 Session ID and removes the whole session context. UPF responds with an N4 session delete response message containing any information that the UPF has to provide to the SMF.
4	SMF interacts with the network entity which triggered this procedure. For example, AMF or PCF.

N4 Session/Node Level Reporting Procedure Call Flows

Session Level Reporting Due to the GTP-u Error Indication Call Flow

When the UPF receives the GTP-u Error Indication from gNB, it detects the PFCP session and sends the PFCP Session Report request to the SMF handling that session along with the Error Indication Report IE. The Error Indication IE also includes the remote F-TEID IE, which contains the GTP-u peer address and the TEID received from the GTP-u Error Indication IE.

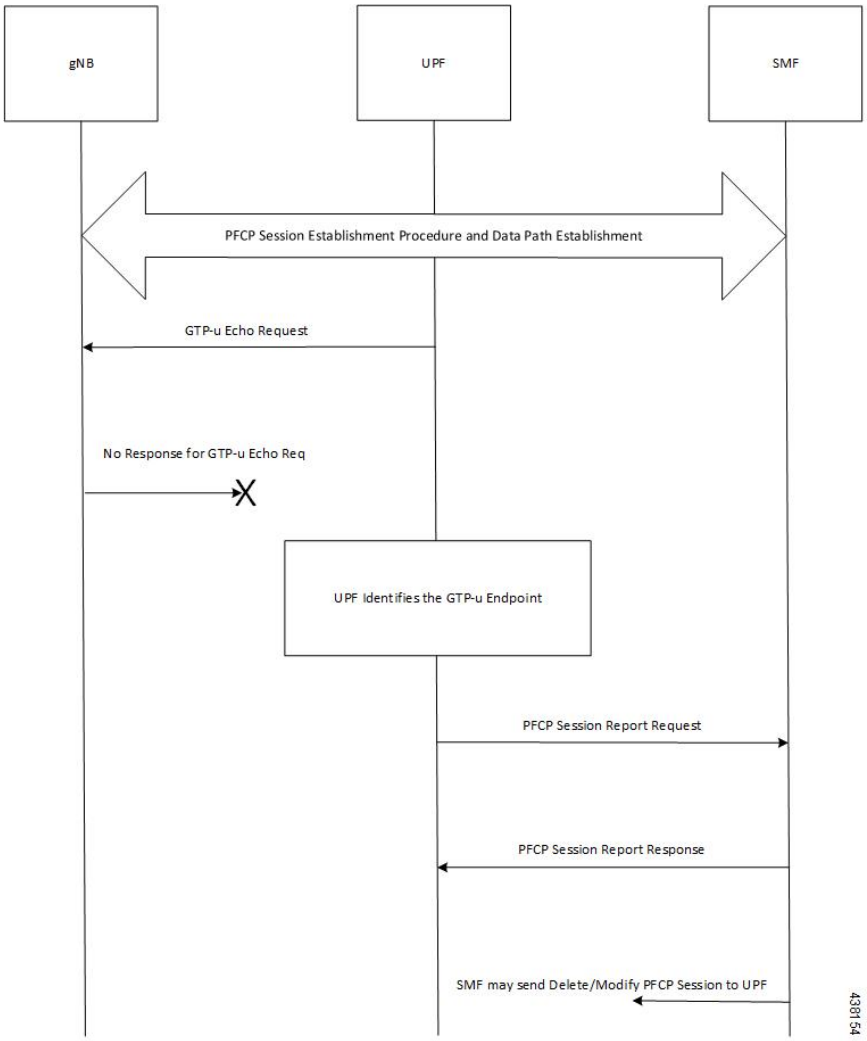


Step	Description
1	A PFCP session is established and the data traffic starts running.
2	When gNB clears up the TEID details locally for some reason, it sends the GTP-u Error Indication to UPF for unknown TEID.
3	Once the Error Indication is received, the UPF identifies the PFCP session and sends the PFCP session report request to the SMF.
4	The session report request contains the TEID and the remote IP address from where the Error Indication is received.

Node-level Reporting Procedure due to GTP-u Path Failure Call Flow

When the UPF enables GTP-u Echo procedure for GTP-u endpoints and identifies a data path failure because of no response, it sends a PFCP Node Report Request to the SMF. The Node Report Type in the PFCP Node Report Request is set to User Plane Path Failure Report when it is sent to the SMF. The Node Report procedure

includes only the peer IP address in Remote GTP-u Peer IE – the child IE of the User Plane Path Failure IE – since it is not specific to any PFCP session.



Step	Description
1	Once the PFCP session is established and GTP-u Echo procedure is configured, UPF initiates the GTP-u Echo Request for each peer with at least one GTP-u Tunnel.
2	If there is no GTP-u Echo Response received after a specified number of retries, then the UPF sends the PFCP Node Report Request to the SMF.
3	Only the peer IP address is sent in the Node Report request since it is not a GTP-u tunnel-specific failure.
4	Once the message is received, the SMF sends a Delete and Modify request for all the PFCP Sessions for that gNB to UPF.

PDN Update Procedure - eNodeB F-TEIDu

Feature Description

For S-GW or SAEGW, a procedure to initiate an N4/Sx Modification Request is implemented for:

- eNodeB F-TEIDu update
- Release Access Bearer (RAB) Request for an eNodeB release

How it Works

The PDN update procedure includes the following events for an eNodeB F-TEIDu Update/Release:

- For eNodeB F-TEIDu Update:
  1. The SGW-C initiates N4/Sx Session Modification Request toward SGW-U on receiving a Modify Bearer Request for eNodeB F-TEIDu Update from the MME.
  2. The N4/Sx Modification Request for eNodeB F-TEIDu update contains Update FAR with Apply Action as “Forward” and the updated eNodeB IPv4/IPv6 address in Outer Header Creation, which is a part of the Update Forwarding Parameters IE.
- For eNodeB F-TEIDu Release:
  1. The SGW-C initiates N4/Sx Modification Request toward SGW-U on receiving a RAB Request from the MME.
  2. RAB is a UE-level message. If the UE has multiple PDN connections, then the N4/Sx Modification Request is sent to each PDN connection separately.
  3. SGW-C initiates N4/Sx Session Modification Request toward SGW-U for the N4/Sx session with Update FAR with destination interface as ACCESS. Update FAR contains: FAR ID and Apply Action as Drop. FAR with the destination interface as CORE is not updated.

Standards Compliance

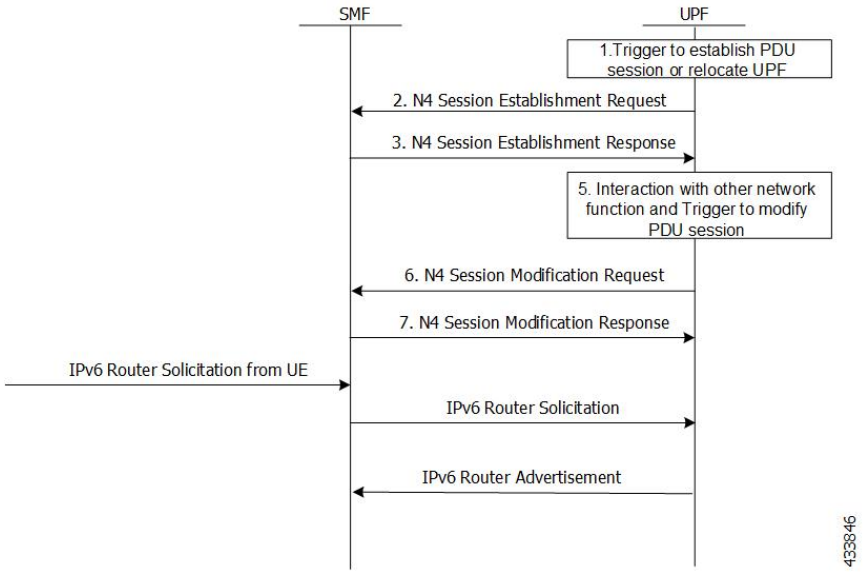
The PDN Update procedure complies with the following standards:

- 3GPP TS 23.401: "General Packet Radio Service (GPRS) enhancements for Evolved Universal Terrestrial Radio Access Network (E-UTRAN) access".
- 3GPP TS 29.274: "3GPP Evolved Packet System (EPS); Evolved General Packet Radio Service (GPRS) Tunnelling Protocol for control plane (GTPv2-C); Stage 3".
- 3GPP TS 29.244: “Interface between the Control Plane and the User Plane of EPC ".
- 3GPP TS 23.214: "Architecture enhancements for control and user plane separation of EPC nodes; Stage 2"
- 3GPP TS 23.714: "Study on control and user plane separation of EPC nodes"

UEs IPv4, IPv6, and IPv4v6 Support Call Flows

N4 Session Establishment and Modification Procedure for IPv6 Call Flow

The following call flow provides a high-level description of the N4 session establishment and modification procedure for IPv6.

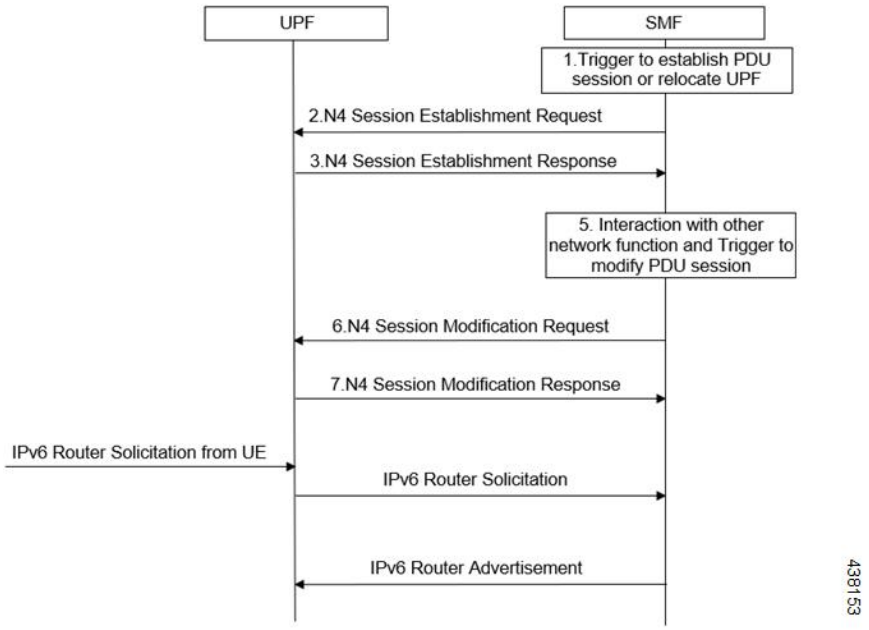


Step	Description
1	The SMF receives the trigger to establish a new PDU session or change the UPF for an established PDU session.
2	The SMF sends N4 Session Establishment Request message to the UPF, in which Create PDR IE has IPv6 UE address, and SDF filter has IPv6 filters to select the PDRs.
3	The UPF responds with an N4 Session Establishment Response message containing any information that the UPF must provide to the SMF in response to the control information received.
4	The SMF interacts with the network function and triggers to modify the PDU session.
5	The SMF sends an N4 Session Modification Request message to the UPF that contains the update for the structured control information which defines how the UPF needs to behave.
6	The UPF identifies the N4 session context to be modified by the N4 Session ID. Then, the UPF updates the parameters of this N4 session context according to the list of parameters sent by the SMF. The UPF responds with an N4 session modification response message containing any information that the UPF must provide to the SMF in response to the control information received.
7	An extra procedure is required for IPv6 sessions, which is Router Solicitation (RS) and Router Advertisement (RA). After the session is established, UE sends the RS message to network to get the link layer address. The UPF forwards this message to SMF. The SMF sends RA message with required parameters to configure the IPv6 address of UE.

Step	Description
8	The RS and RA between UPF and SMF is GTP-u encapsulated and SMF sends an extra pair of PDRs during session establishment and modification procedure, in which GTP-u Tunnel IDs are exchanged for GTP-u tunneling.
9	The additional pair of PDRs that are sent from the SMF are as follows: • One PDR has Source interface as Access, and Destination Interface as CP-Function, to forward IPv6 RS from UE to SMF. The SDF filter is present so that UPF can select this PDR for IPv6 RS from UE • Example: Permit in 58 from any tp ff01::2 any • Another PDR has source interface as CP-Function, and Destination Interface as Access, to forward the IPv6 RA from SMF to UE. The SDF filter is present so that UPF can select this PDR for all the IPv6 RAs from SMF. • Example: Permit out 58 from any to ff01::2 any
10	After RS and RA procedure is completed, the UE sends IPv6 traffic to PDN.

N4 Session Establishment and Modification Procedure for IPv4v6 Call Flow

The following call flow provides a high-level description of the N4 session establishment and modification procedure for IPv4v6.



The IPv4v6 session establishment and modification procedure are similar to the IPv6 session establishment and modification procedure, except for the following procedures:

Step	Description
1	In the session establishment request, a PDR IE is created to include both IPv4 and IPv6 UE addresses. The SDF filter also includes the IPv4 and IPv6 filters for selecting the PDRs.
2	When the IPv6 address is assigned to the UE, an extra procedure – Router Solicitation and Router Advertisement – is required. The UE sends the Router Solicitation message to the network to receive the link layer address, once the session is established. The UPF forwards this message to the SMF and the SMF sends the Router Advertisement Message with the required parameters for configuring the IPv6 address of the UE.
3	The RS/RA between UPF and SMF is GTP-u encapsulated. In addition, the SMF sends an extra pair of PDRs during session establishment and modification procedure, in which the GTP-u tunnel IDs are exchanged for GTP-u tunneling.
4	The additional pair of PDRs that are sent from the SMF are as follows: • For forwarding IPv6 Router Solicitation from UE to SMF, one PDR's source interface is set to access, and its destination interface is set to CP-Function. The SDF filter is present in such a way that the UPF selects this specific PDR for the IPv6 Router Solicitation from the UE. • For instance, permit in 58 from any to ff01::2 any.
5	The UE sends the IPv6 traffic to PDN once the RS/RA procedure is complete.
6	No additional procedures are required for the IPv4 traffic for this PFCP session.

Configuring the N4 Session/Node Level Reporting Procedures

This section describes how to configure the N4 Session/Node Level Reporting procedures.

Enabling the GTP-u Echo Request Procedure

The existing CLI (Command Line Interface) in **gtpu-service** is used to enable the GTP-u Echo request procedure.

```
configure
   gtpu-service service_name
      echo-interval seconds
      echo-retransmission-timeout seconds
      max-retransmissions num
      path-failure detection-policy gtp echo
      end
```

NOTES:

- **gtpu-service** *service_name*: Creates a GTP-u service enters the GTP-u Service Configuration Mode for the current context. *service_name* specifies the name of the GTP-u service.

- **echo-interval** *seconds*: Configures the rate at which GTP v1-u echo packets are sent. *seconds* specifies the number of seconds between the sending of a GTP-uv1 echo packet. It must be an integer in the range of 60–3600.
- **echo-retransmission-timeout** *seconds*: Configures the timeout for GTP-u echo message retransmissions for this service. *seconds* specifies the echo retransmission timeout, in seconds, for the GTP-u service. It must be an integer in the range of 1–20. The default value is 5.
- **max-retransmissions** *num*: Configures the maximum retry limit for GTP-u echo retransmissions. *num* specifies the number of GTP-u echo message retransmissions allowed before triggering a path failure error condition. It must be an integer in the range of 0–15.
- **path-failure detection-policy gtp echo**: Configures a path failure detection policy on GTP-u echo messages that have been retransmitted the maximum number of retry times**. gtp echo** sets the detection policy to detect a failure upon reaching the maximum number of GTP-u echo message retransmissions.

The following is a sample configuration for enabling GTP-u Echo request procedure.

```
configure
   gtpu-service n3-gtpu-service
   echo-interval 60
   echo-retransmission-timeout 5
   max-retransmissions 5
   path-failure detection-policy gtp echo
   end
```

Verifying the N4 Session/Node Level Reporting Procedure Configuration

This section describes how to verify the N4 Session/Node Level Reporting Procedure configuration.

N4 Session Node Level Reporting Procedure OA and M Support

Use the **show gtpu statistics command** to display the GTP-u statistics for Error Indication and GTP-u Echo Request and Response. The following is a sample output from the **show gtpu statistics** command.

```
show gtpu statistics
Path Management Messages:
    Echo Request Rx:                  0 Echo Response Rx:            7
    Echo Request Tx:                 19 Echo Response Tx:            0
    SuppExtnHdr Tx:                  0  SuppExtnHdr   Rx:            0

  Peer Stats:
    Total GTPU Peers:                              1
    Total GTPU Peers with Stats:                   1

  Tunnel Management Messages:
    Error Indication Tx:                           0
    Error Indication Rx:                           0
    Error Indication Rx Discarded:
```

Use the **show sx-service statistics all** command to display the Node report request and response statistics. The following is a sample output of the **show sx-service statistics all** command.

```
show sx-service statistics all
Node Report Request:
 Total TX:                          1   Total RX:                      0
 Initial TX:                        1   Initial RX:                    0
 Retrans TX:                        0   Retrans RX:                    0
 No Rsp received TX:                0   Discarded:                     0
```

```
Node Report Response:
 Total TX:                          0   Total RX:                          1
 Initial TX:                        0   Initial RX:                        1
 Accepted:                          0   Accepted:                          1
 Denied:                            0   Denied:                            0
 Retrans TX:                        0   Discarded:                         0
```

Use the **show user-plane-service statistics all** command to display the statistics of N4 PFCP message parameters. The following is a sample output of the **show user-plane-service statistics all** command.

```
show user-plane-service statistics all
N4 Statistics:
    URR : Created        :              0
          Deleted        :              0
          Queried by ID  :              0
          Queried by all :              0
          Total Queried  :              0
    FAR : Created        :              0
          Updated        :              0
          Removed        :              0
    PDR : Predef rule    :              0
```

Note Statistics are cumulative, and are displayed as a total of all session managers.

The descriptions of the fields are as follows:

- URR Created: Displays the total number of URRs created either locally or as requested by the SMF.
- URR Deleted: Displays the total number of URRs removed either locally or as requested by the SMF.
- URR Queried by ID: Displays the total messages received to query a subset of URRs with their specific IDs.
- URR Queried by all: Displays the total messages received to query all the URRs of a session.
- URR Total Queried: Displays the total number of URRs reported to the SMF in response to "Queried by ID" and "Queried by all".
- FAR Created: Displays the total number of FARs created.
- FAR Updated: Displays the total number of FARs updated.
- FAR Removed: Displays the total number of FARs removed.
- PDR Predef rule: Displays the total number of CREATE PDRs that have predefined rules set. For one predefined rule two PDRs are received, and so the counter is updated as 2.

Note Total Predef counter isn't updated when a predefined rule is removed. SMF only sends Remove PDR message, which doesn't contain predefined rule name. Also, during HO, for exiting predef rule, PDRs are removed and created again, so it's counted twice for the same session.

SNMP Traps

The following traps are available to track status and conditions GTP-u path failure.

- EGTPUPathFailure: This trap is generated when no response is received for GTP-U ECHO requests and data path failure is detected toward a peer EPC Node.
- EGTPUPathFailureClear: This trap is generated when the data path toward the peer node is available.

Bulk Statistics

The following bulk statistics counters are newly added:

- Identify type of change (new/modified)
- Schema name
- Statistic (Variable) Name
- Key Variable (yes/no)
- Description
- Triggers
- Availability
- Stat Type
- Standard or Proprietary
- Data Type

Enhanced PFCP Association Release Procedure for Graceful Session Termination

Table 83: Feature History

Feature	Release	Description
PFCP Association Release Request Initiation from UPF	2024.03.0	This feature lets UPF send notification on PFCP Session Release to cnSGWc and SMF. The UPF notification indicates to clear the calls simultaneously in UPF and SMF, or UPF and cnSGWc. If the SMF or cnSGWc is not notified, the call remains connected until UPF receives the next Session Modify Request from the SMF or cnSGWc. This leads to loss of subscriber usage reports. Here, the Enhanced PFCP Association Release (EPFAR) feature improves the signaling efficiency and effective handling of usage reports by SMF or cnSGWc. **Default Setting:** Disabled – Configuration Required to Enable

When the UPF decides to clear the call due to an error or a partial failure of a session, the UPF clears the calls locally without informing the SMF or cnSGWc on the call clearance. The call remains connected until the next Session Modify Request received by UPF from the SMF or cnSGWc.

To avoid losing the usage reports and improving the signaling efficiency during call clearance, the Enhanced PFCP Association Release (EPFAR) feature is applied for the UPF to initiate the session report request to gracefully clear the session between UPF and SMF or UPF and cnSGWc simultaneously. This feature complies with the Release 16.9.0 of 3GPP TS 29.244, section 5.18.1 and, section 5.18.2.

EPFAR Negotiation

EPFAR feature negotiation is an aggregate communication between the UPF and SMF or UPF and cnSGWc. When UPF and SMF or UPF and cnSGWc support the EPFAR feature, the session and association release actions are accomplished. You can enable the EPFAR feature using the configuration command when the UPF needs to release the association with the SMF and cnSGWc.

The association release process involves these procedures:

1. UPF-initiated PFCP Session Release
2. UPF-initiated Enhanced PFCP Association Release

UPF-initiated PFCP Session Release

UPF-initiated PFCP Session Release process is executed through these steps:

1. UPF enables the EPFAR feature for a peer node only if it is negotiated during Association Setup procedure. When the UPF needs to delete a PFCP sesssion due to an error or a partial failure, it initiates the PFCP Session Report Requests for the affected session. UPF sends PFCP Session Report Request to SMF or cnSGWc with Report Type and Usage Report Trigger IEs.
   - The Report Type is set as USAR (Usage Report) when there is a non-zero usage report for the PFCP session or UISR (UP Initiated Session Request) if there is no usage report to send.

     The fifth bit of Octet 6 in Report Type IE is a proprietary IE bit used for indicating UISR for PFCP Session Report Request.
   - Usage Report Trigger is set as TEBUR (Termination By UP function Report) for a non-zero usage report.
2. UPF sends the Cause IE to the peer node in PFCP Session Report Request message. The Cause IE uses these values for communicating the cause of session deletion:
   - 201 - Subscriber Clear
   - 202 - Association Release initiated by UP
   - 203 - Recovery Failure
   - 204 - IP Source Violation
   - 205 - PFCP Cause Self Protection Termination
3. UPF sets the PSDBU (PFCP Session Deleted By the UP function) flag as 1 to indicate the PFCP session deletion.
4. UPF receives the PFCP Session Report Response from the SMF or cnSGWc and deletes the sessions locally in UPF and SMF or UPF and cnSGWc.

Note The SMF has a clear rate limitation of 500 sessions per second. Hence, the session report throttle rate is maintained at 500 sessions per peer interface at UPF.

Clear the Subscriber Session in UPF

To clear the subscriber sessions in UPF, enter the **pace-out-interval** value in the **clear subscribers all** command.

Procedure

Step 1 Use the formula to calculate the **pace-out-interval** value.

Example:

```
pace_per_sessmgr = throttle rate / num_of_smgr
pace-out-interval = max session in one smmgr / pace_per_sessmgr

Throttle Rate = 500 sessions / second
Subscriber data-rate = 203000
Total number of smgrs = 22
pace_per_sessmgr = throttle rate / num_of_smgr = 500 / 22 = 22
pace-out-interval = maximum sessions in one smmgr / pace_per_sessmgr = (203000 / 22) / 22  = 420
```

Step 2 Enter the **pace-out-interval** value to the command.

Example:

```
[local]upf# clear subscribers all pace-out-interval pace-out-interval

[local]upf# clear subscribers all pace-out-interval 420
```

Recovery Handling in UPF

The EPFAR feature configuration in both Active and Standby chassis need to be same. The configured EPFAR feature is not checkpointed. If the chassis moves from Standby to Active state, the Standby chassis reads the information from configuration.

Note If the graceful cleanup of dropped calls is enabled using the **require graceful-cleanup-during-audit-failure** CLI command, the UPF will not initiate the PFCP Session Report Request message for the sessions during recovery failure.

Verify the Cause Level Statistics

Use the **show sx-service statistics all** command to verify cause level information from the the node level statistics.

Procedure

Enter the show command to display the cause level information.

Example:

```
[local]qvpc-si# show sx-service statistics all

Session Management Messages:
...
...
...
Session Report Request:
 Total TX:                              350   Total RX:                    0
 Initial TX:                            0   Initial RX:                  0
 Retrans TX:                            0  Retrans RX:                   0
 Discarded:                             0  No Rsp RX:                    0
 Causes:
     User Initiated Session Deletion from UP  36
     Association Release initiated by UP      310
```

```
Recovery Failure                        0
IP Source Violation                     4
```

UPF-initiated Enhanced PFCP Association Release

1. When both the SMF or cnSGWc and UPF support EPFAR feature, UPF initiates the PFCP Association Report Request with PARPS (PFCP Association Release Preparation Start) flag set from the UPF on PFCP Association Release and the SMF or cnSGWc stops selecting the UPF.
2. UPF initiates the PFCP Session Release procedure and after all reports are sent, UPF sends PFCP Association Update Request and set the URSS (non-zero Usage Reports) flag to 1 for the affected PFCP Sessions Sent.
3. On reception of successful Update response from the SMF or cnSGWc, UPF starts the configurable graceful cleanup timer to release the association. You can enable the release timer using the configuration command to set the timer. The SMF or cnSGWc will initiate the deletion of affected sessions.
4. If the UPF receives PFCP Association Release Request from SMF or cnSGWc, it stops the timer and releases the association. Else, UPF releases the association after the timer expiration.

Note The session report throttle rate is maintained at 500 sessions per peer interface at UPF in Association Release procedure.

PFCP Association Release without EPFAR

When the peer node does not support EPFAR feature or the EPFAR is not enabled on UPF, PFCP Association Release is executed through these steps:

1. The PFCP association release is initiated in UPF using the configuration command.
2. UPF sends the PFCP Association Update Request with SARR flag set in PFCP Association Release Request IE.
3. UPF starts the timer after receiving update response from the SMF or cnSGWc. The SMF or cnSGWc deletes the affected sessions.
4. If the timer expires, UPF deletes all the affected PFCP sessions and PFCP association locally.

Enable the EPFAR Feature

Enable the EPFAR feature on UPF using the **sx-protocol supported-features epfar** command.

Procedure

Step 1 Enter the Context Configuration mode for an UP.

Example:

[local]UPF1(config)# **context EPC2-UP**

Step 2 Enter the *service_name* in Sx Service mode.

Example:

[EPC2-UP]UPF1(config-ctx)# **sx-service name** *service_name*

Step 3 Enable EPFAR on UPF to intiate the session or association release. By default, the EPFAR feature is in disabled state.

Example:

[EPC2-UP]UPF1(config-sx-service)# **sx-protocol supported-features epfar**

Use the **no sx-protocol supported-features epfar** command to disable EPFAR.

Step 4 Save the configuration.

Step 5 Enter the show command to verify if the EPFAR feature is enabled or disabled.

Example:

```
[local]qvpc-si# show sx peers full address 192.0.2.0
  Peer IP                     : 192.0.2.0
  Sx Service Id               : 5
  Group Name                  : cp1
  Current Session             : 0
  Max Session                 : 0
  Negotiated feature
  Load Control                : Disabled
  Overload Control            : Disabled
  EPFAR                 : Enabled
```

Initiate the Association Release for Sx Service

The **clear sx-association peer-node-id ipv4-address** command initiates the association release for Sx association with an IPv4 or IPv6 peer-node-id in UPF.

Procedure

Enter the Sx association for peer-node-id and mention the IP address for an IPv4 or IPv6 peer.

Example:

```
[local]UPF1# clear sx-association peer-node-id ipv4-address
 IPv4 Address
```

```
[local]UPF1#clear sx-association peer-node-id ipv6-address
 IPv6 Address
```

```
clear sx-association peer-node-id ipv4-address 192.0.2.0
```

Enable the Association Release Timer

The **sx-protocol association release-timeout** command allows to enable the association release timer in UPF.

Procedure

Step 1 Enter the Context Configuration mode command for an UP.

Example:

```
[local]UPF1(config)# context EPC2-UP
```

Step 2 Enter the *service_name* and *release_timer* in Sx Service mode. You can set the timer between 30 to 900 seconds. The default release-timeout value is 90 seconds.

Example:

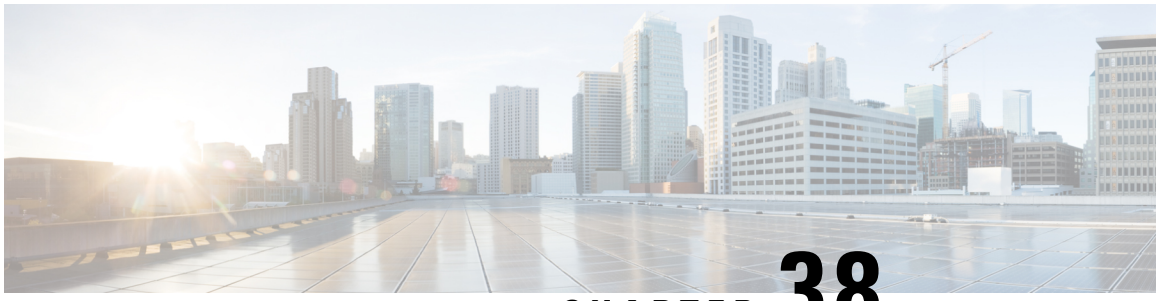
```
[EPC2-UP]UPF1(config-ctx)# sx-service epc
[EPC2-UP]UPF1(config-sx-service)# sx-protocol association release-timeout 100
```

Step 3 Enter the show command to verify the Association Release Timeout and EPFAR feature information for Sx service.

Example:

```
[local]qvpc-si# show sx-service name sxu

  Service name                          : sxu
  Service-Id                            : 4
  Context                               : ingress
  ...
  SX Association Reattempt Timeout      : 900 (sec)
  SX Association Release Timeout        : 30 (sec)
  ...
  SX PDI Optimisation                   : Enabled
  Supported Features :
       EPFAR                            : Enabled
```



CHAPTER 38

NRF Support

- Feature Summary and Revision History, on page 407
- Feature Description, on page 408
- How it Works, on page 409
- Configuring NRF Management Services, on page 409
- Monitoring and Troubleshooting, on page 411

Feature Summary and Revision History

Summary Data

Table 84: Summary Data

Applicable Product(s) or Functional Area	5G-UPF
Applicable Platform(s)	VPC-SI SMI
Feature Default Setting	Disabled - Configuration Required
Related Changes in this Release	Not Applicable
Related Documentation	Not Applicable

Revision History

Table 85: Revision History

Revision Details	Release
First introduced.	2022.04.0

Feature Description

In the 5G service-based architecture, the Network Repository Function (NRF) maintains an updated repository of all the 5G Network Functions (NFs) available in the operator's network. NRF also contains the details of the services provided by the 5G NFs, and allows the 5G NFs to instantiate, scale, and terminate without or minimal manual intervention.

NRF interacts with all NFs in the 5G core network, and provides the following services:

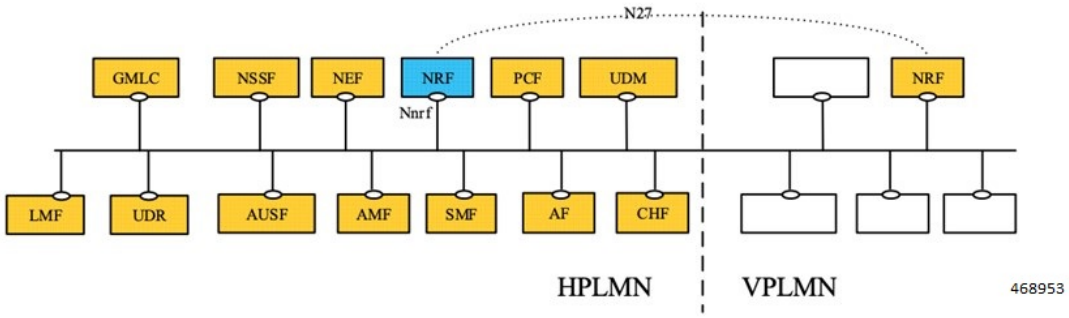
- Management Services
- Discovery Services
- OAuth2 Authorization
- Bootstrapping

The UPF supports only NRF Management Services.

NRF Management Services

The NRF Management (Nnrf_NFManagement) service enables the NF instances in the serving PLMN to register, update, or deregister their profiles in the NRF. The NF instance described here is the UPF StarOS based instance.

Figure 25: NRF Management



Presently, the Nnrf_NFManagement service provides the following operations:

- **Register NF instance** (NFRegister)—An NF instance registers its profile in the NRF along with the list of services that it provides.
- **Update NF instance** (NFUpdate)—An NF instance partially updates or replaces the NF profile parameters in the NRF. It also adds or deletes the services that it provides.

  Currently, NFUpdate supports the following operations:

  - Update of NF profile attributes
  - Heartbeat of NF

- **De-register NF instance** (NFDeregister)—An NF instance deregisters its NF profile and the services that it provides in the 5G network.

How it Works

This feature enables management services between UPF and NRF.

UPF Registration

Once the minimal configuration is done, the UPF triggers the registration procedure toward the NRF. The Nrfmgr proclet receives the configuration from the SessCtrl, and selects the endpoint with the lowest priority-number. If Nrfmgr does not receive response from the NRF for a retry, then Nrfmgr selects the endpoint with the next lowest priority-number. This process continues with all the endpoints endlessly. To stop this process, you must delete the configuration.

UPF Heartbeat

Once the registration is complete, Nrfmgr checks if the "nr hb-enable" option is configured. If "nr hb-enable" is configured already, Nrfmgr builds or sends the Heartbeat Request message. Nrfmgr then starts the timer using the heartbeat timer-value received from peer NRF. If no heartbeat response is received until the heartbeat timer-value, Nrfmgr picks the endpoint with the next lowest priority-number. This process continues with all the endpoints in a round-robin fashion endlessly. To stop this process, you must delete the configuration.

UPF DeRegistration

Once the configuration is deleted, UPF sends an NRF Deregistration message to NRF from Nrfmgr, and thus NRF deregisters UPF.

Standards Compliance

The NRF Support feature complies with *3GPP TS 29.510 "5G System; Network function repository services; Stage 3"*.

Configuring NRF Management Services

NRF Profile Configuration

The NRF Profile configuration provides UPF the flexibility or control to define the optional parameters that must be sent toward the NRF. For example, if you configure locality in the NNRF Mgmt profile, and the NF profile is associated with the NNRF service, locality is sent to the NRF.

To configure the NRF profile in the UPF, use the following CLI commands:

```
configure
  context context_name
    user-plane-nnrf
      nnrf-mgmt-format profile_name
        priority number
        locality string
```

```
      sst sst_number [ sd sd_number ]
      smf-serving-area area_n
    exit
  exit
```

NOTES:

- **user-plane-nnrf**: Configures the UPF NRF profile.
- **nnrf-mgmt-format** *profile_name*: Configures the NNRF management profile to control the parameters that are sent to NRF.
- **priority** *number*: Specifies the priority of the UPF. If configured, the value is sent to NRF.

  number must be an from integer from 1 to 100. Default: No value is sent to NRF.
- **locality** *string*: Specifies the locality of the UPF. If configured, this value is sent to NRF. The maximum length of *string* must be 63 bytes.
- **sst** *sst_number* **[sd** *sd_number* **]**: Specifies the Slice information of the UPF. If configured, this value is sent to NRF.

  sst *sst_number* must be an integer from 0 to 255. Default: Slice/SST is not sent to NRF.

  sd *sd_number* must be an integer from 0 to 16777215. If not configured, SD is not sent to NRF.
- **smf-serving-area** *area_n*: Specifies the SMF serving area of the UPF. You can configure up to five SMF areas.

  The maximum length of *area_n* must be 63 bytes. If configured, the value is sent to NRF.

NRF Service Configuration

To configure NRF services in the UPF, use the following CLI commands:

```
configure
  context context_name
    nnrf-nfm-service service_name
      associate nnrf-mgmt-format profile_name
      uri-scheme { http | https }
      hb-enable
      retransmission-timeout timeout_value
      max-retransmissions max_retries
      certificate path [ key path ] [ ca-certificate path ]
      endpoint-name endpoint_name
        priority number
        ipv4-address ipv4_address [ portv4 port_number ] [ ipv6-address
ipv6_address ] [ portv6 port_number ]
        exit
      bind ipv4-address ipv4_address [ portv4 port_number ] [ ipv6-address
ipv6_address ] [ portv6 port_number ]
    exit
    user-plane-service userplane_service
      associate nnrf-nfm-service service_name
      nf-instance-id number
    exit
```

```
  exit
exit
```

NOTES:

- **nnrf-nfm-service** *service_name*: Configures the NRF service.
- **associate nnrf-mgmt-format** *profile_name*: Associates the NNRF management profile configuration. If not associated, the profile configurations such as *priority/slice/smfarea* are not sent in NFMgmt messages to the NRF.
- **uri-scheme { http | https }**: Specifies the URI Scheme that is used to send messages to NRF (HTTP or HTTPS). Default value: HTTP.
- **[no] hb-enable**: Enables or disables heartbeat messages to NRF.
- **retransmission-timeout** *timeout_value*: Specifies the retry timer-interval, in seconds, for an endpoint to send messages to NRF. *timeout_value* must be an integer from 1 to 100. Default value: 15 seconds.
- **max-retransmissions** *max_retries*: Specifies the maximum retries for sending messages to NRF. *max_retries* must be an integer from 1 to 20. Default value: three retries.
- **certificate** *path* **[key** *path* **] [ca-certificate** *path* **]**: Specifies the path to certificate or CA certificate, and key URLs.
- **endpoint-name** *endpoint_name*: Specifies the NRF endpoint. You must configure at least one endpoint to trigger messages from UPF to NRF.
- **priority** *number*: Specifies the NRF endpoint priority.

  number must be an integer from 1 to 100. Default value: 50.
- **bind ipv4-address** *ipv4_address* **[portv4** *port_number* **] [ipv6-address** *ipv6_address* **] [portv6** *port_number* **]**: Binds an IPv4 or IPv6 address to NRF. This configuration is required to trigger messages from UPF to NRF. If configured, IPv6 is prioritized over IPv4.
- **user-plane-service** *userplane_service*: Defines the **user-plane-service** to **nnrf-nfm-service** association and the NF instance. This configuration is required to trigger messages from UPF to NRF.
- **associate nnrf-nfm-service** *service_name*: Specifies the associated **nnrf-nfm-service** service.
- **nf-instance-id** *number*: Specifies the NF instance ID in UUID format.

Monitoring and Troubleshooting

This section provides information about the CLI commands available for monitoring and troubleshooting this feature.

Show Commands and/or Outputs

This section describes the clear and show CLI commands for this feature.

- **clear nrf statistics**
- **show nrf statistics**

- **show nrf nnrf-nfm-service all**
- **show nrf nnrf-nfm-service name** *service_name*

  The following is a sample output of this command:

```
  Service name:             nrf-svc1
 VPN Name / Id:             ingress / 2
  State       :             Started
  Timer Value :             15
  Max Retries :             3
  Heartbeat   :             Enabled
 Bind IPv4Addr:             209.165.201.2
 Bind IPv4Port:             0
 Bind IPv6Port:             0
   URI Scheme :             HTTPS
   NRF profile:             prof1
```

- **show nrf nnrf-nfm-service name** *service_name* **statistics**

  The following is a sample output of this command:

```
Service  name:             nrf-svc1
  Curr Endpoint:             end1
  Curr State   :             NRFMGMT_STATE_PENDING_REG

Statistics for endpoint name end1
=====================================================================

Num PUT Req      : 3                    Num PATCH Req           : 0

Num PUT Success  : 0                    Num PATCH Success       : 0

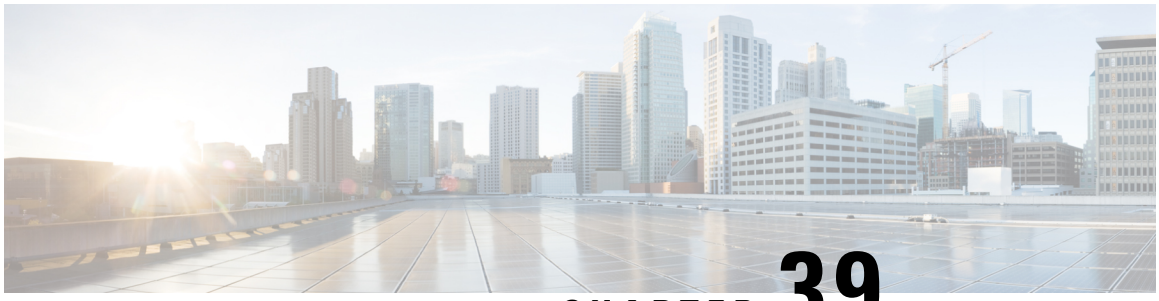
Num PUT Failed   : 3                    Num PATCH Failed        : 0

Num DEL Req      : 0               Last Conn Req Time      : 2022-08-05+19:11:14

Num DEL Success  : 0                    Last Conn Resp Time     : NA

Num DEL Failed   : 0
```

- **mon-pro support** *18*



CHAPTER 39

Password Expiration Notification

- Feature Summary and Revision History, on page 413
- Feature Description, on page 414
- Upgrading and Downgrading Procedures using Save Configuration Command, on page 415

Feature Summary and Revision History

Summary Data

Table 86: Summary Data

Applicable Product(s) or Functional Area	5G-UPF
Applicable Platform(s)	VPC-SI SMI
Feature Default Setting	Disabled – Configuration Required
Related Changes in This Release:	Not Applicable
Related Documentation	*UCC 5G UPF Configuration and Administration Guide*

Revision History

Table 87: Revision History

Revision Details	Release
First introduced.	2022.01.0

Feature Description

If the password isn't reset before the expiration date, you get locked from the UPF. You're allowed to log in back only when the password is reset by the administrators manually.

UPF provides password expiration notification to Context/AAA/Radius users. UPF supports configuration and expiration of passwords for Administrators, Config Administrators, Inspectors, and Operators. Following provisions are supported:

- Specify the password warning interval. It warns you about password expiry.
- Specify the password grace interval. During this grace interval, you can change the password by yourself rather than approaching the Administrator every time.
- Warning interval and Grace interval have a global configuration under a context. If the user level configuration doesn't specify either of these values, global values under the context take effect.

The default values of the parameters are according to the Security Guidelines.

- Expiry Interval—Maximum age of the password (default: 90 days)
- Warn Interval—Warning period before password expiry (default: 30 days). You get a warning about approaching password expiry. You can continue without changing the password.
- Grace Interval—Days after password expiry you can use the old password. Beyond the grace period, you may not be able to log in with the old password. Admin has to reset the password for you.

For example:

```
login: xxx
password: xxx

Case 1: [Normal]
# {you are logged in}

Case 2: [When in warning period]
Warning: Your password is about to expire in 0 days.
We recommend you to change password after login.
Logins are not allowed without acknowleding this.
Do you wish to continue [y/n] (times out in 30 seconds) :

Case 3: [when in grace period]
Your password has expired
Current password:
New password:
Repeat new password:

Case 4: [after the grace period]
Password Expired (even beyond grace period, if configured). Contact Security Administrator
to reset password
```

Upgrade and Downgrade Enhancement for Password Expiration Notification

Password Expiry Notification feature has introduced many new keywords in Subscriber configuration such as **max-age**, **exp-grace-interval**, and **exp-warn-interval**. These new parameters are configured at the Context Global level. Context Global level parameters are used when per user level configuration isn't configured with a default value. For example, for the **max-age** of the password, the default value is 90 days.

For the user profiles with no "expiry-date" at per user level, startup config takes an expiry date of 90 days for that user. This problem can be solved by manually editing the startup configuration file, but this solution leads to issues when users are distributed across locations.

In case if downgrade is needed, user profiles are lost as new keywords aren't valid for older releases.

With the password expiration notification enhancement, the upgrade procedure is updated, and the downgrade process is changed with the help of new **save config** CLI option, **legacy-password-expiry**.

Upgrading and Downgrading Procedures using Save Configuration Command

Use the following upgrade process:

- Before upgrade, add **no password max-age** command at context level, in all contexts where users are configured, in the startup configuration.
- When reloading with image using the updated startup config, all users that are configured without an expiry date pick up the context level configuration by default and set the user level **no-max-age** keyword automatically.

Use the following downgrade process:

Use the new CLI option, **legacy-password-expiry** in the save config command, based on which new keywords aren't saved. Configuration is stored in a format which the previous release recognizes.

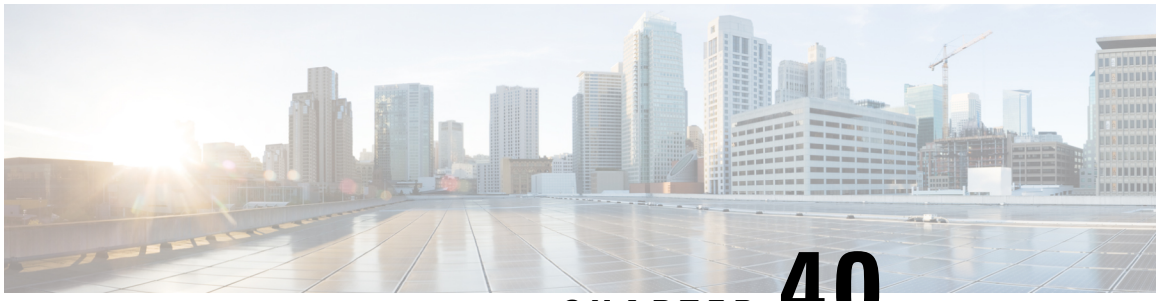
The following prompt is displayed in the Exec mode:

```
configure
   context host_name
      save configuration url [ confd | ignore-locks | obsolete-encryption
| showsecrets | verbose ] [ -redundant ] [ -noconfirm ] [
legacy-password-expiry ]
```

Notes:

- **save configuration** *config-file-path* **legacy-password-expiry**:

  Generates a backward compatible file by removing new Expiry Notification keywords. The **save config** CLI option makes the configuration compatible with older UPF versions.



CHAPTER 40

QCI 80 Support on UPF

- Feature Summary and Revision History, on page 417
- Feature Description, on page 418
- How it Works, on page 418
- Configuring ADC Rule, on page 420
- Monitoring and Troubleshooting, on page 421

Feature Summary and Revision History

Summary Data

Table 88: Summary Data

Applicable Product(s) or Functional Area	5G-UPF
Applicable Platform(s)	VPC-SI SMI
Feature Default Setting	Enabled – Always-on
Related Changes in this Release	Not Applicable
Related Documentation	*UCC 5G UPF Configuration and Administration Guide*

Revision History

Table 89: Revision History

Revision Details	Release
First introduced.	2022.01.1

Feature Description

The 5G-UPF supports new standard QoS Class Index (QCI) 80 based on 3GPP TS 23.203, for establishing a non-GBR QoS flow when an application sends traffic to the specific destination.

How it Works

Dynamic QoS Flow Establishment based on Detected Traffic

To establish a dynamic QoS flow when traffic is detected, UPF uses the Application Detection and Control (ADC) over Gx feature.

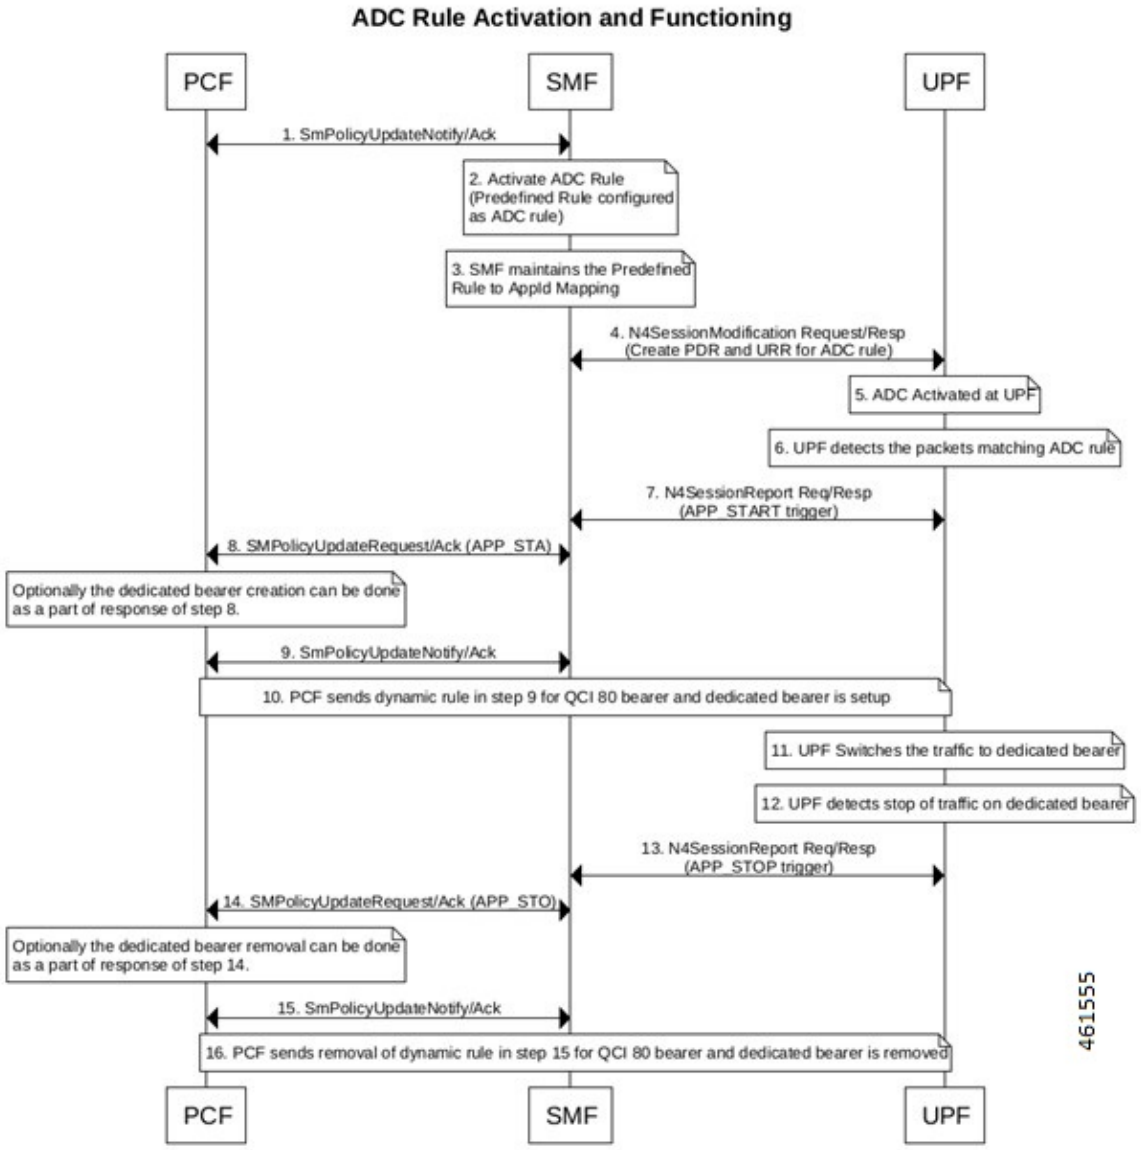
On receiving a PCC predefined rule over default 5G QoS Identifier (5QI) for application detection and control, the SMF instructs the UPF to detect application traffic. The UPF installs the Gx ADC PDR with default QFI, and the rule name is considered as Application ID for reporting to the SMF.

On detecting a new flow hitting the Gx ADC predefined rule, the UPF triggers an Application Start event. When the application traffic is identified by an application identifier, the SMF reports the start of application to the PCF.

The PCF then makes policy decisions based on the information received and sends the corresponding updated PCC rules (with QCI-80) to the SMF.

Call Flow

The following figure shows the call flow details and the message exchanges in a 5G core system.



Step	Description
1	The PCF sends SmPolicyUpdateNotify to SMF, and receives an acknowledgment.
2	SMF activates the ADC rule (Predefined rule).
3	SMF maintains the predefined rule for Application ID mapping.
4	SMF sends the N4 Session Modification Request to UPF, and receives a response from UPF.
5	ADC is activated at the UPF.
6	UPF detects the packets matching ADC rule.

Step	Description
7	UPF sends the N4SessionUpdate request to the SMF, and receives a response from the SMF. UPF triggers Application Start.
8	SMF sends the SmPolicyUpdate Request to PCF, and receives an acknowledgment from the PCF. Optionally, the dedicated bearer is created as part of response.
9	PCF sends SmPolicyUpdateNotify to SMF, and receives an acknowledgment from SMF.
10	PCF sends dynamic rule for QCI 80 bearer, and a dedicated bearer is setup.
11	UPF switches the traffic to dedicated bearer.
12	UPF detects stop of traffic on dedicated bearer.
13	UPF sends the N4SessionUpdate request to the SMF, and receives a response from the SMF. UPF triggers Application Stop.
14	SMF sends the SmPolicyUpdate Request to PCF, and receives an acknowledgment from the PCF. Optionally, the dedicated bearer is removed as part of response.
15	PCF sends SmPolicyUpdateNotify to SMF, and receives an acknowledgment from SMF.
16	PCF sends removal of dynamic rule for QCI 80 bearer, and the dedicated bearer is removed.

Limitations

The following is the known limitation to this feature in this release:

- Gx ADC predefined rule is installed only with default bearer 5QI.

Configuring ADC Rule

To support activation of predefined ADC rules, you must configure the ADC rule in SMF with appropriate action priority, and **adc** keyword. The following is a sample configuration:

```
active-charging service service_name
    rulebase rulebase_name
          action priority action_priority dynamic-only adc ruledef ruledef_name
 charging-action charging_action_name
    exit
ruledef ruledef_name
   ip server-ip-address ipv4/ipv6_address/mask
   ip server-ip-address ipv4/ipv6_address/mask
   exit
exit
```

To mute the reporting, use the following CLI command under rulebase configuration:

action priority *action_priority* **dynamic-only adc mute ruledef** *ruledef_name* **charging-action** *charging_action_name*

To optimize the application reporting once per application, use the following CLI command under rulebase configuration:

adc app-notification once-per-app

To optimize the application reporting once per application per flow, use the following CLI command under rulebase configuration:

adc app-notification once-per-ipflow

NOTES:

- **rulebase** *rulebase_name*: Enables the Active Charging Service Rulebase configuration
- **action priority** *action_priority*: Assigns priority to a ruledef in the rulebase. Priority must be a unique integer value ranging 1–65535.
- **dynamic-only**: Enables matching of dynamic rules with static rules for this action priority on a flow.
- **adc**: Specifies the ruledef to-be given as ADC rule.
- **ruledef** *ruledef_name* : Adds the specified ruledef to the current rulebase.
- **charging-action** *charging_action_name*: Specifies the charging action.
- **description** *description*: Adds specified text to the rule and action.
- **ip server-ip-address** *ipv4/ipv6_address/mask*: Specifies the server IP address with subnet mask bit. For uplink packets (subscriber to network), this field matches the destination IP address in the IP header. For downlink packets (network to subscriber), this field matches the source IP address in the IP header.
- **once-per-app**: Notifies APP_START or APP_STOP notification once per App ID.

  Whenever traffic for an application is detected, the UPF sends APP-START notification to the SMF with an Application Instance Identifier as the Flow-Id of the first data flow of an application.This application instance Id gets cached in the User Plane.When the last data flow of an application is terminated, the cached application instance Id is sent in APP-STOP notification so that PCF can better correlate APP-START and APP-STOP.
- **once-per-ipflow**: Notifies APP_START or APP_STOP notifications per App ID per IP flow.

Monitoring and Troubleshooting

This section provides information on CLI commands that are available for monitoring and troubleshooting the feature.

Show Commands to Display PDR and URR

Use the following command on UPF to display the PDR for which ADC is enabled:

show subscribers user-plane-only callid *callid_value* **pdr full all**

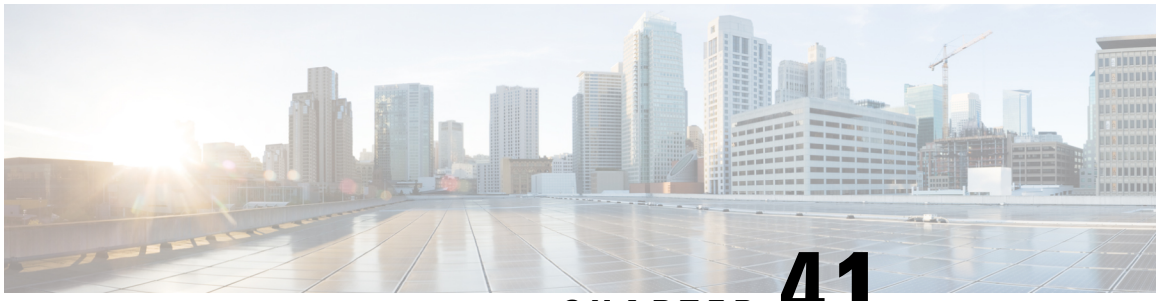
Use the following command on UPF to display the URR for which ADC is enabled:

show subscribers user-plane-only callid *callid_value* **urr full all**

Show Commands to Display ADC Statistics

Use the following command on UPF to display the ADC statistics:

show subscribers user-plane-only callid *callid_value* **adc statistics**



CHAPTER 41

QoS Group of Ruledefs Support

- Feature Summary and Revision History, on page 423
- Feature Description, on page 424
- How it Works, on page 424
- OAM Support, on page 426

Feature Summary and Revision History

Summary Data

Table 90: Summary Data

Applicable Product (s) or Functional Area	5G-UPF
Applicable Platforms	VPC-SI SMI
Feature Default Setting	Disabled – Configuration Required
Related Changes in this Release	Not Applicable
Related Documentation	*UCC 5G UPF Configuration and Administration Guide*

Revision History

Revision Details	Release
First introduced.	2023.01.0

Feature Description

The QoS Group of Ruledefs feature helps in enforcing fair usage policy (FUP) per subscriber. QoS Group of Ruledefs is also referred to as QGR or SGQ.

The QGR feature sets different QoS parameters for different subscribers for a named QGR, therefore ensuring fair usage policing for a subscriber.

How it Works

The QGR feature performs flow status and bandwidth limiting under the charging-action configuration. UPF applies the static configuration for QGR using RCM.

When a packet matches a ruledef, UPF performs a QGR match to check if a QGR is present with the matched ruledef or group. UPF returns the highest priority QGR.

- If a QGR matches, UPF applies flow-action enforcement at charging-action, and then at QGR if charging-action allows the packet. If the packet drops, UPF skips the QGR-level flow-action enforcement.
- If flow-action at QGR allows the packet, UPF enforces the QoS Enforcement Rule (QER) limiting on a packet. If the packet drops at QGR, UPF skips QER limiting.
- UPF performs QER limiting at charging-action and then at QGR if the packet is allowed at charging-action.

UPF supports QoS-group-of-ruledefs configuration under the ACS service.

The following is a sample QGR configuration under ACS:

```
configure
   active-charging service acs
     qos-group-of-ruledefs QGR1
        add-group-of-ruledef group
        add-ruledef http
        #exit
     rulebase test
        action priority 2 ruledef http charging-action standard
        action priority 5 ruledef catchall charging-action standard
        route priority 1 ruledef http-rule analyzer http
        end
```

The following configuration is an example of QoS Group "QGR1" received over PCRF:

```
-qos-group-rule-install
--qgr-name QGR1
--qgr-mon-key 1
--qgr-flow-status 2
--qgr-precedence 1
--qgr-eqos-information
--qgr-eqos-mbr 1000 2000
--qgr-eqos-mbr-burst-size 1000 2000
--qgr-eqos-mbr-limit-conform-action 0 -1 0 -1
--qgr-eqos-mbr-limit-exceed-action 1 -1 1 -1
```

IE Support

UPF pushes the QGR in Session Establishment Request or Session Modification Request using the private IE—QGR-INFO-LIST. If there are changes to the QGR dynamic parameters, UPF triggers an Update FAR or QER. UPF sends the updated value in Session Establishment Request or Session Modification Request.

Currently, UPF does not support the monitoring-key parameter in QGR. Hence, UPF does not send URR associated with QGR in Session Establishment Request or Session Modification Request.

The following is the name and format of the private IE:

```
QGR-INFO-LIST:
Value:
Number of QGRs: 2
QGR INFO:
NAME: qgr-gor
PRECEDENCE: 3
OPERATION: QGR ADD
FAR ID: 0x8005
QER ID: 0x0004
QGR INFO:
NAME: qgr2
PRECEDENCE: 9
OPERATION: QGR ADD
FAR ID: 0x8006
QER ID: 0x0005
```

The following table describes the FAR format:

Table 91: FAR Format

FAR ID	Unique ID
Apply Action	Flow-action defined in QGR.
Extended Apply Action	Private IE to communicate to QGR FAR. This IE includes Flow-Action—Discard Uplink, Discard Downlink, and Terminate Flow.
Forwarding Parameters	Contains the destination interface and redirect information if any.

The following table describes the QER format:

Table 92: QER Format

QER ID	Unique ID

Maximum Bitrate	The maximum bit rate of QGR in Kbps.
	`UL MBR`
	`DL MBR`

Burst Size	Private IE to configure the burst size.
	`UL Burst`
	`DL Burst`

Conform Action	Private IE to configure the conform action.
	`Uplink ToS`
	`Downlink ToS`

QER ID	Unique ID

Exceed Action	Private IE to configure the exceed action.
	`Uplink Action`
	`Uplink ToS`
	`Downlink Action`
	`Downlink ToS`

Use the **show subscribers user-plane-only callid** *call_id* **{ far | qer } full all** command to view the FAR and QER values.

Limitations

The QGR feature has the following limitations:

- The monitoring key associated with QGR is not usage monitored, that is, UPF does not support URR creation and enforcement.
- UPF does not support the inclusion of dynamic rules in the static QGR definition.
- UPF does not support flow-Status redirect and kill flow.
- UPF does not support QoS Group Conform action as Drop and Exceed action as ALLOW or MARK_DSCP.
- UPF supports a maximum configuration of 64 QGRs statically.
- UPF supports a maximum of 64 QGRs received from SMF.

OAM Support

This section describes the operations, administration, and maintenance information for this feature.

Show Commands Support

This section provides information regarding show commands and their outputs in support of the QGR feature.

show subscribers user-plane-only callid call_id far full all

This command displays the Extended Apply Action IE information on UPF.

```
[local]qvpc-si# show subscribers user-plane-only callid 00004e21 far full all
FAR-ID                   : 0x8004
         Destination Interface    : Core
         Apply Action             : FORWARD
         Outer Header Creation    :
         Remote TEID              : 0x0
         Transport Level Marking         : N/A
         Transport Level Marking Options :
           Copy Inner:                     No
           Copy Outer:                     No
         Inner Packet Marking            : N/A
         Remote IP Address        :
         Remote Port              : N/A
```

```
Number of Associated PDRs: 0
Associated with QGR      : Yes
    Extended Apply Action : ULDROP DLDROP
```

show subscribers user-plane-only callid call_id qer full all

This command displays the bandwidth parameters associated with QGR on UPF.

```
[local]qvpc-si# show subscribers user-plane-only callid 00004e21 qer full all

        QER-id                  : 0x00000003        Correlation-id          : n/a
        GBR uplink (bps)        : 0                 GBR downlink (bps)      : 0

        UL Gate Status          : OPEN              DL Gate Status          : OPEN
        Number of Associated PDRs  :0
        Associated with QGR     : Yes
        MBR uplink (bps)        : 1000              MBR downlink (bps)      : 2000
        UL Burst         : 1000                    DL Burst              : 1000

           UL Conform Action          : ALLOW                 DL Conform
Action     : ALLOW
           UL DSCP Value              : NA                    DL DSCP Value
         : NA
            UL Exceed Action          : DROP                DL Exceed Action
      : DROP
           UL DSCP Value              : NA                    DL DSCP Value
         : NA
```

show subscribers user-plane-only callid call_id qos-group all

This command displays the QGR information from QGR-INFO-LIST IE on UPF.

```
[local]qvpc-si# show subscribers user-plane-only callid call_id qos-group
all       name   statistics
[local]qvpc-si# show subscribers user-plane-only callid 00004e21 qos-group all

Callid: 00004e21
Interface Type: Sxb
    QGR-Name:      Priority:       FAR-ID:         QER-ID:         URR-ID:
    ---------      ---------       -------         -------         ------
     QGR1           5               0x8003          0x0002          0x0000
     QGR2           1               0x8004          0x0003          0x0000

Total Number of QGRs found: 2
```

show subscribers user-plane-only callid call_id qos-group statistics

This command displays QoS-group statistics on UPF.

```
[local]qvpc-si# show subscribers user-plane-only callid 00004e21 qos-group statistics
all       name      |         <cr>
QoS-Group-of-Ruledef:
    Total Uplink Pkts:                0     Total Dnlink Pkts:                0
    Total Uplink Bytes:               0     Total Dnlink Bytes:               0
    Uplink Pkts Allowed:              0     Dnlink Pkts Allowed:              0
    Uplink Bytes Allowed:             0     Dnlink Bytes Allowed:             0
    Uplink FP Pkts:                   0     Dnlink FP Pkts:                   0
    Uplink FP Bytes:                  0     Dnlink FP Bytes:                  0
     Flow-Status Statistics:
       Total Uplink Packets:           0     Total Downlink Packets:           0
```

```
      Total Uplink Bytes:             0     Total Downlink Bytes:             0
      Uplink Packets Redirected:      0     Downlink Packets Redirected:      0
      Uplink Bytes Redirected:        0     Downlink Bytes Redirected:        0
      Uplink Packets Dropped:         0     Downlink Packets Dropped:         0
      Uplink Bytes Dropped:           0     Downlink Bytes Dropped:           0
      Uplink Packets Term-Flow:       0     Downlink Packets Term-Flow:       0
      Uplink Bytess Term-Flow:        0     Downlink Bytes Term-Flow:         0
    Bandwidth-Control Statistics:
      Total Uplink Packets:           0     Total Downlink Packets:           0
      Total Uplink Bytes:             0     Total Downlink Bytes:             0
      Uplink Packets QoS-Exceed:      0     Downlink Packets QoS-Exceed:      0
      Uplink Bytes QoS-Exceed:        0     Downlink Bytes QoS-Exceed:        0
      Uplink Packets QoS-Conform:     0     Downlink Packets QoS-Conform:     0
      Uplink Bytes QoS-Conform:       0     Downlink Bytes QoS-Conform:       0
      Uplink Packets Dropped:         0     Downlink Packets Dropped:         0
      Uplink Bytes Dropped:           0     Downlink Bytes Dropped:           0
      Uplink Packets Marked:          0     Downlink Packets Marked:          0
      Uplink Bytes Marked:            0     Downlink Bytes Marked:            0

Total qos-group-of-ruledefs matched: 2

Total subscribers matching specified criteria: 1
```

show subscribers user-plane-only full all

This command displays the number of per subscriber QoS-groups that UPF has successfully activated and the QoS-group statistics matched for data packets.

```
[local]qvpc-si# show subscribers user-plane-only full all

  Local SEID      : [0x0004000000000000] 1125899906842624
  Remote SEID     : [0x0004000000000000] 1125899906842624
  State           : Connected
  Connect Time    : Wed Mar 18 07:53:57 2020
  Idle time       : 18h53m24s
  Access Type: uplane-ipv4              Network Type: IP
  user-plane-service-name: user-plane-service
  Callid: 00004e21
  Rulebase: prepaid
  Interface Type: Sxb
  Card/Cpu: 1/0                         Sessmgr Instance: 1
  IP address: 10.0.0.1
  Source context: ingress               Destination context: egress
  PDN-Instance: intershat
  User-plane-Sx-addr: 192.60.182.1
  Control-plane-Sx-addr: 192.60.181.1
  Number of associated PDRs : 0
  Number of associated ADC PDRs : 0
  Number of associated FARs : 1
  Number of associated QERs : 2
  Number of associated URRs : 3
  ….
  ….
  ….
  CC Dropped Uplink bytes: 0                      CC Dropped Downlink Bytes: 0
  Uplink Inflight Pkts: 0                         Downlink Inflight Pkts: 0
  Total QoS-Group(s) Active: 1

  Flow information:
    Current Active Flows:
      TCP: 0
      UDP: 0
    Total Flows:
```

```
      TCP: 2
      UDP: 0
      FP:  0

Static & Predef Rule Match stats:
Rule Name            Pkts-Down  Bytes-Down Pkts-Up    Bytes-Up   Hits       Match-Bypassed
  FP-Down(Pkts/Bytes)  FP-Up(Pkts/Bytes)
-------------------- ---------- ---------- ---------- ---------- ---------- --------------
  -------------------  -----------------
http                          2       1352          2        250          4              0
          0/0              0/0
catchall                      4        160          8        320         12              0
          0/0              0/0

Dynamic Rule Match stats:
PDR Id   Pkts-Down  Bytes-Down Pkts-Up    Bytes-Up   Hits       Match-Bypassed
FP-Down(Pkts/Bytes)    FP-Up(Pkts/Bytes)
-------- ---------- ---------- ---------- ---------- ---------- --------------
-------------------    -----------------

Post-Processing Rule Match stats:
Rule Name            Pkts-Down  Bytes-Down Pkts-Up    Bytes-Up   Hits       Match-Bypassed
  FP-Down(Pkts/Bytes)  FP-Up(Pkts/Bytes)
-------------------- ---------- ---------- ---------- ---------- ---------- --------------
  -------------------  -----------------
QoS-Group Statistics:
QGR Name  Pkts-Down   Bytes-Down    Pkts-Up    Bytes-Up    Hits
Match-Bypassed   FP-Down(Pkts/Bytes)  FP-Up(Pkts/Bytes)
-------- ----------   ----------   ---------  ----------  ------
--------------   -------------------  -----------------
QGR1        3            756           5            285          8         0
               0/0                    0/0
```

show user-plane-service qos-group-of-ruledefs

This command displays the static QoS-groups configuration received on UPF through PFD.

```
[local]qvpc-si# show user-plane-service qos-group-of-ruledefs
all       name
[local]qvpc-si# show user-plane-service qos-group-of-ruledefs all
  QOS-Group-of-Ruledefs Name: QGR1
   Ruledef Name
   ============
   group

Total QoS group(s) of ruledefs found: 1
```

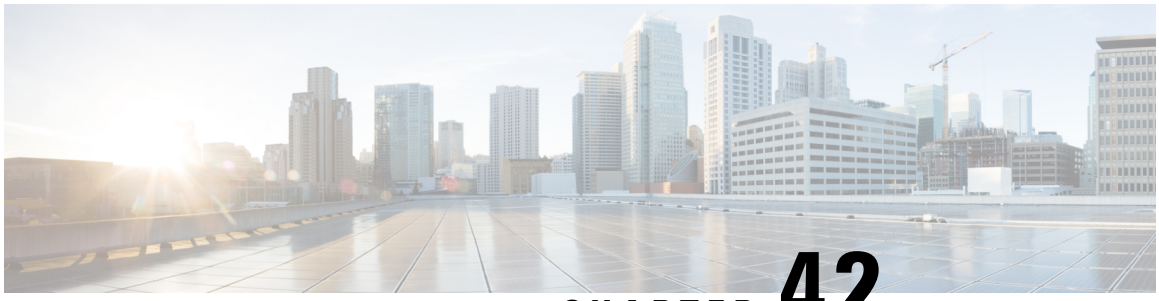
show user-plane-service statistics qos-group sessmgr

This command displays cumulative statistics across all QoS-Groups.

```
[local]laas-si-setup# show user-plane-service statistics qos-group sessmgr
1..1152   all

Sessmgr Instance: 1
      Total Uplink Pkts:                5       Total Dnlink Pkts:                3
      Total Uplink Bytes:               274     Total Dnlink Bytes:               756
      Uplink FP Pkts:                   0       Dnlink FP Pkts:                   0
      Uplink FP Bytes:                  0       Dnlink FP Bytes:                  0
      Flow-Status Statistics:
        Total Uplink Packets:           5       Total Downlink Packets:           3
        Total Uplink Bytes:             274     Total Downlink Bytes:             756
```

```
  Uplink Packets Redirected:       0       Downlink Packets Redirected:        0
  Uplink Bytes Redirected:         0       Downlink Bytes Redirected:          0
  Uplink Packets Dropped:          0       Downlink Packets Dropped:           0
  Uplink Bytes Dropped:            0       Downlink Bytes Dropped:             0
  Uplink Packets Term-Flow:        0       Downlink Packets Term-Flow:         0
  Uplink Bytess Term-Flow:         0       Downlink Bytes Term-Flow:           0
Bandwidth-Control Statistics:
  Total Uplink Packets:            5       Total Downlink Packets:             3
  Total Uplink Bytes:            274       Total Downlink Bytes:             756
  Uplink Packets QoS-Exceed:       0       Downlink Packets QoS-Exceed:        0
  Uplink Bytes QoS-Exceed:         0       Downlink Bytes QoS-Exceed:          0
  Uplink Packets QoS-Conform:      5       Downlink Packets QoS-Conform:       3
  Uplink Bytes QoS-Conform:      274       Downlink Bytes QoS-Conform:       756
  Uplink Packets Dropped:          0       Downlink Packets Dropped:           0
  Uplink Bytes Dropped:            0       Downlink Bytes Dropped:             0
  Uplink Packets Marked:           0       Downlink Packets Marked:            0
  Uplink Bytes Marked:             0       Downlink Bytes Marked:              0
```



CHAPTER 42

Roaming Support

- Feature Summary and Revision History, on page 431
- Feature Description, on page 432
- How it Works, on page 435
- Configuring the HR Roaming Support for UPF, on page 452
- Monitoring and Troubleshooting, on page 454

Feature Summary and Revision History

Summary Data

Table 93: Summary Data

Applicable Product(s) or Functional Area	5G-UPF
Applicable Platform(s)	VPC-SI SMI
Feature Default Setting	Enabled – Always-on
Related Changes in this Release	Not Applicable
Related Documentation	Not Applicable

Revision History

Table 94: Revision History

Revision Details	Release
UPF supports inter-PLMN roaming mobility and secondary PDR for 5G sessions.	2023.03.0

Revision Details	Release
UPF supports the following roaming functionalities: • Source Interface Type IE indicates the 3GPP Interface Type • Subscriber Params IE indicates the roaming status	2023.02.0
First introduced.	2021.02.0

Feature Description

The mobile network operators form roaming partnerships to provide seamless services to their subscribers in geographies beyond their network reach. Operator network boundaries are designated by public land mobile networks (PLMN). The home network for a subscriber is called an HPLMN and the visited network, which renders the mobile service is termed as the VPLMN.

The VPLMN provides access network services and packet routing to the packet core, whereas the HPLMN provides data network access to the subscriber. This feature enables the UPF to support the flavor of routing that is termed as the Home Routed (HR) roaming.

Visited UPF (vUPF) supports the following functionalities:

- Handle the dummy PDRs with associated FAR action as buffer
- Buffer the incoming packets before rule matching
- Handle the QoS and FAR updates from the SMF for debuffering of packets
- Send the buffered packets after matching with the PDR
- Send vUPF traffic over Fast Path
- Send N9 interface GTP-U tunnel
- UPF LI support is extended for roaming
- MonSub CLI command and PCAP file support

vUPF and Home UPF (hUPF) support the following functionalities:

- QoS flow Based Charging (QBC) on the UPF
- Source Interface Type IE in PDR indicates the 3GPP Interface Type in UPF
- Selection of the GTP-U service by receiving 3GPP Interface Type received in PDI IE of create PDR.
- Subscriber Params IE indicates the roaming status in UPF

Inter-PLMN Mobility Support

When UE moves from its home network to a visited network, the Public Land Mobile Network (PLMN) changes from home-PLMN to visited-PLMN. This handover is referred to as inter-Public Land Mobile Network

(PLMN) mobility. UPF supports inter-PLMN with optimization and secondary PDR optimization for 5G and 4G sessions.

For information on the call flows related to inter-PLMN handover, see the *How it Works* section.

Relationships

The roaming status functionality relates to the UPF Ingress Interfaces feature that indicates the following types of roaming:

- Homer
- Roamer
- Visited LBO
- Visited HR

The UPF supports different GTP-U ingress interfaces. It allows a separate network for the N3 interface and N9, S5u, and S8u interfaces.

The GTP-U ingress interfaces to handle roaming and outbound calls include:

- N3—N3 is the interface between gNodeB and UPF. To handle 5G home calls, the N3 interface must be a private IP address.
- N9—The N9 interface connects two UPFs. It is the interface between intermediate I-UPF and UPF session anchor connecting different PLMN. To handle 5G outbound roaming calls, the N9 interface must be a public IP address.
- S5u—S5u is similar to the N9 interface that connects two UPFs. The interface connects the intermediate I-UPF and UPF session anchor. To handle 4G inbound calls, the S5u is set to a public IP since the SGW-U is public for enabling the S-GW.
- S8u—S8u is an inter-PLMN variant of the S5u interface. To handle 4G outbound roaming calls, the S8u interface must be a public IP address.

For more information, refer to the *UPF Ingress Interfaces* chapter.

Architecture

This section describes the architecture for the home routing roaming support feature.

Bufffering and Debuffering on the vUPF

The buffering and debuffering procedure on the vUPF for the UPF HR roaming, are as follows:

1. Two dummy PDRs (UL or DL) is created initially at vUPF by vSMF with default value as QFI and the buffering as an FAR action.
2. The packets coming from the N3 and N9 interfaces get buffered based on the FAR action before doing a packet classification and application of policy.
3. If buffered packet count exceeds the configured limits then the subsequent packets are dropped.

4. The buffered packets are sent for classification and policy application after the update FAR is received with action as forward and updated TEID in modify request from vSMF.
5. The vSMF initiates the removal of default QER and URR while sending the update FAR with action forward and sends a new PDR with the required QER and URR.
6. The packets are sent on the required interface that is based on the QFI defined in the new QER.
7. If there is no matching PDR with TEID and QFI installed then the debuffered packets is dropped.

Charging — Predef and Dynamic Rules on the hUPF

The charging predef and dynamic rules on the hUPF, are as follows:

1. The SMF associates FBC URRs + QBC URRs + Session URRs with dynamic PDRs.
2. The SMF associates QBC URRs + Session URRs with predef PDRs.
3. The UPF associates the URRs created by installed global PDRs to the received predef PDRs QBC URRs + Sess URRs.
4. The QBC URRs have no Linked URRs.
5. The QBC URRs include no FBC URRs or Session URRs usage reporting.
6. The UPF links the Session URRs to FBC URRs as Linked URRs.
7. The Session URRs include FBC URRs and QBC URRs usage reporting.
8. The UPF relies on the SMF for the update or removal of each of the QBC and Sess URRs.

Charging — Static Rules on the hUPF

The charging static rules on the hUPF, are as follows:

1. The SMF associates FBC URRs + QBC URRs + Session URRs with RB PDRs.
2. The UPF associates the URRs created by installed global PDRs to the received RB PDRs QBC URRs + Sess URRs.
3. The UPF does not link QBC URRs with any URRs.
4. The UPF links the static FBC URRs with the Session URRs.
5. The UPF links the Session URRs to the FBC URRs as linked URRs for usage reporting.
6. The UPF relies on the SMF for the update or removal of each of the QBC and the SMF URRs.

Inter-PLMN Mobility Support

The Inter-Public Land Mobile Network (Inter-PLMN) mobility feature allows the PLMN to change from home-PLMN to visited-PLMN when a User Equipment (UE) moves from its home network to a visited network and vice-versa. This feature provides support for various types of Inter-PLMN mobility cases.

For information on the call flows related to inter-PLMN handover, see the *How it Works* section.

H-UPF Data Path specifications:

- Secondary PDR optimization for Homer calls
- Inter-PLMN and Intra-PLMN handover support with optimization
- Supports GTPU services for Homer calls
- KPI support for Secondary PDR info in PDR
- Handover statistics

UPF supports the following scenarios for various Inter-PLMN mobility:

- 5G homer to 5G roamer
- 5G homer to 4G roamer
- 4G homer to 5G roamer
- 4G homer to 4G roamer
- 5G roamer to 5G homer
- 5G roamer to 4G homer
- 4G roamer to 5G homer
- 4G roamer to 4G homer

UPF supports the following scenarios for various Intra-PLMN mobility:

- 5G homer to 4G homer
- 4G homer to 5G homer
- 5G roamer to 4G roamer
- 4G roamer to 5G roamer
- EUTRAN to WiFi and vice-versa handover

How it Works

This section describes the call flows related to the roaming feature in UPF.

PDU Session Establishment Procedure

This call flow describes the PDU Session Establishment procedure for the UPF.

Figure 26: PDU Session Establishment Call Flow

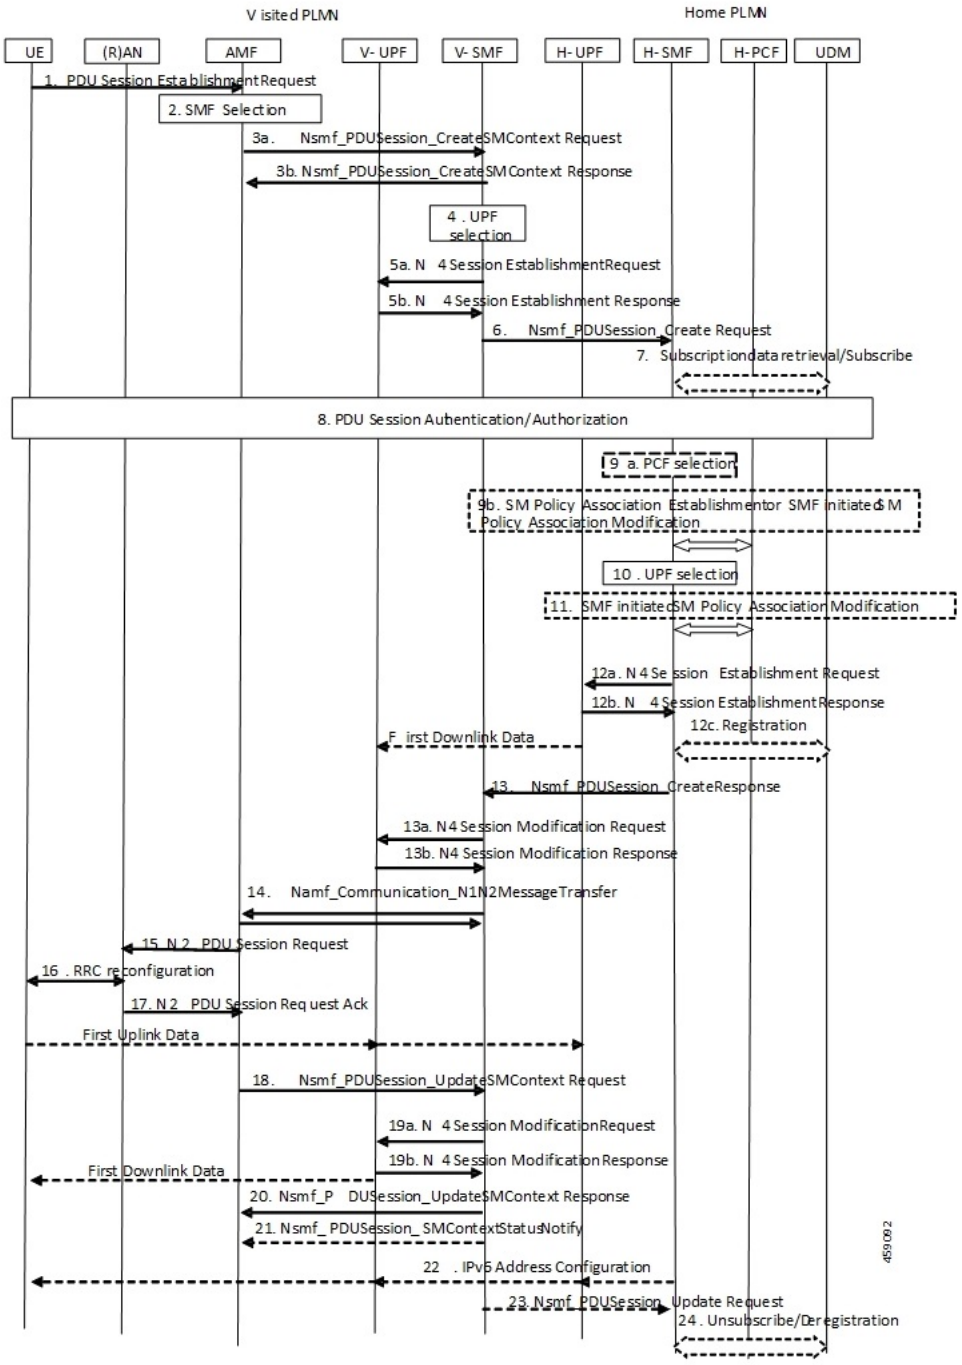


Table 95: PDU Session Establishment Call Flow Description

Step	Description
1	The UE initiates a PDU Session Establishment Request to the AMF.

Step	Description
2	The AMF selects an SMF.
3a	The AMF invokes the Nsmf_PDUSession_CreateSMContext Request and sends it to the vSMF.
3b	The vSMF sends a Nsmf_PDUSession_CreateSMContext Response to the AMF.
4	The vSMF selects a UPF in VPLMN.
5a	The vSMF sends an N4 Session Establishment Request with dummy PDRs, FARs, QERs, or URRs to the vUPF. In this request, the vSMF informs the vUPF to allocate the CN Tunnel information for the N3 and N9 interfaces. The vSMF sends Source Interface Type IE in Access PDR as N3 3GPP Access (11), and in Core PDR as N9 (21) for roaming.
5b	The vUPF creates the N3 and N9 CN tunnel information and acknowledges by sending this information in an N4 Session Establishment Response.
6	The vSMF sends an Nsmf_PDUSession_Create Request to the hSMF.
7	The hSMF registers with the UDM for a given PDU session.
8	The SMF performs a secondary authorization or authentication during the establishment of the PDU session by a DN-AAA server.
9a	The hSMF selects the PCF.
9b	The SMF performs an SM Policy Association Establishment procedure to establish an SM Policy Association with the PCF and get the default PCC rules for the PDU session.
10	The hSMF selects the UPF.
11	The hSMF initiates an SM Policy Association Modification procedure.
12a	The hSMF initiates an N4 Session Establishment procedure with the selected UPF. The N9 CN tunnel information from the vUPF is transferred to hUPF in FAR. The hSMF sends Source Interface Type in Access PDR as N9(21) or N3 3GPP ACCESS (11) or S5u interface depending on the RAT and roaming status where it registers the call.
12b	The UPF acknowledges by sending an N4 Session Establishment Response. The hUPF provides the N9 CN tunnel information to the vUPF.
12c	The downlink path is established between the hUPF and vUPF and data packets are sent to the vUPF where it gets buffered.
13	The hSMF sends an Nsmf_PDUSession_Create Response message to the vSMF.
14	The vSMF sends an Namf_Communication_N1N2MessageTransfer message to the AMF.
15	The AMF sends the NAS message containing PDU Session ID and PDU Session Establishment Accept targeted to the UE and the N2 SM information received from the SMF within the N2 PDU Session Request to the RAN.
16	The RAN issues AN specific signalling exchange with the UE that is related with the information received from SMF.

Step	Description
17	The RAN sends an N2 PDU Session Response message to the AMF.
18	The AMF sends the Nsmf_PDUSession_UpdateSMContext Request message to the vSMF. The AMF forwards the N2 SM information received from RAN to the vSMF.
19a	The vSMF initiates an N4 Session Modification procedure with the vUPF. The vSMF provides packet detection, enforcement, and reporting rules to be installed on the vUPF for this PDU session, including AN Tunnel Information, H-CN Tunnel Information and V-CN Tunnel Information.
19b	The vUPF provides an N4 Session Modification Response to the vSMF. After this step, the vUPF delivers any down-link packets to the UE that might have been buffered for this PDU Session.
20	The vSMF sends an Nsmf_PDUSession_UpdateSMContext Response message to the AMF.
21	The vSMF sends an Nsmf_PDUSession_SMContextStatusNotify message to the AMF.
22	If it's a PDU session of type IPv6 or IPv4v6, the hSMF generates an IPv6 Router Advertisement and sends it to the UE through the N4 interface, hUPF, and vUPF.
23	If the vSMF received in step 18 is an indication that the RAN has rejected some QFIs the vSMF notifies the hSMF through an Nsmf_PDUSession_Update Request. The hSMF is responsible for updating the QoS rules and QoS Flow level QoS parameters for the QoS Flow(s) associated with the QoS rule(s) in the UE accordingly.
24	If the PDU Session establishment failed after step 4, the hSMF performs the following steps: • If the SMF is no more handling a PDU Session of the UE for this (DNN, S-NSSAI), the hSMF unsubscribes to the modifications of Session Management Subscription data for the corresponding (SUPI, DNN, S-NSSAI), by using Nudm_SDM_Unsubscribe. • The hSMF deregisters for the given PDU Session by using Nudm_UECM_Deregistration (SUPI, DNN, PDU Session ID).

PDU Session Modification Procedure

This call flow describes the PDU Session Modification procedure for the UPF.

Figure 27: PDU Session Modification Call Flow

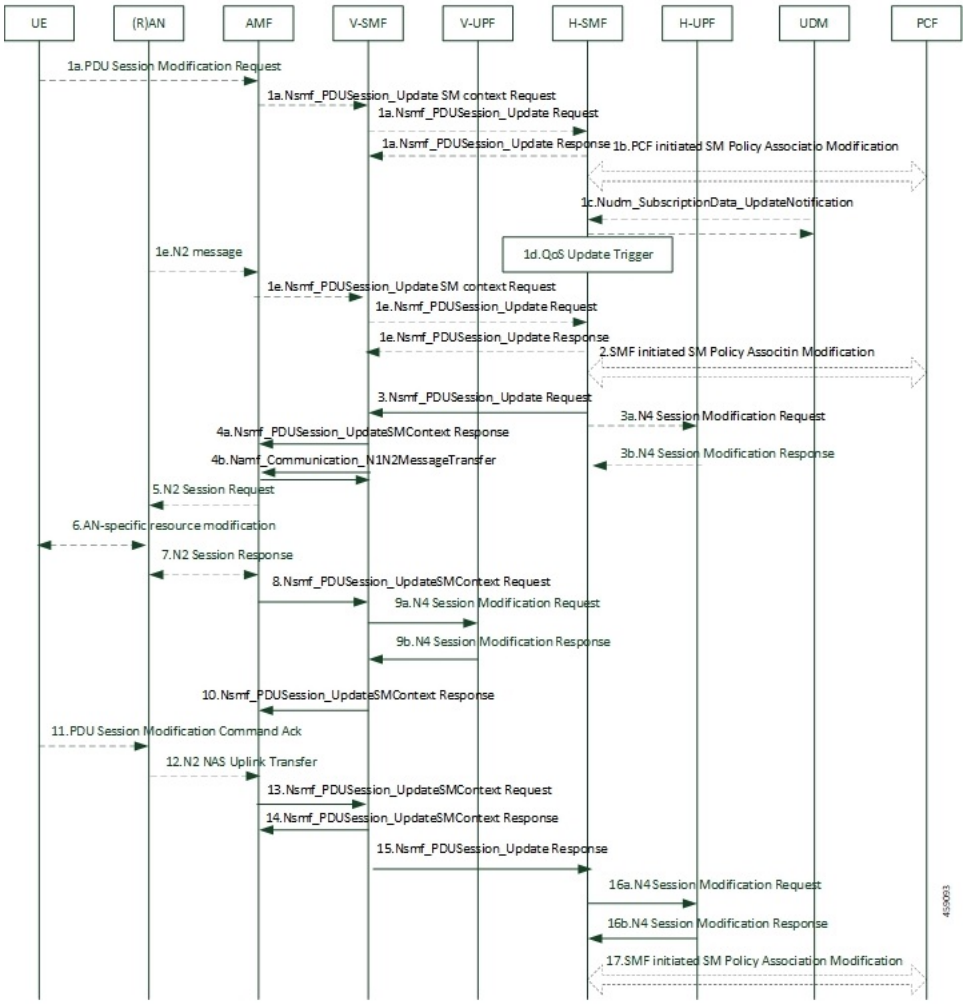


Table 96: PDU Session Modification Call Flow Description

Step	Description
1a	• The UE initiates the PDU Session Modification procedure by the transmission of an NAS message to the AMF. • The AMF initiates the Nsmf_PDUSession_UpdateSMContext message. • The vSMF sends an Nsmf_PDUSession_Update Request message to the hSMF. • The hSMF acknowledges and sends an Nsmf_PDUSession_Update Response message to the vSMF.
1b	The PCF performs a PCF initiated SM Policy Association Modification procedure to notify SMF about the modification of policies.

Step	Description
1c	The UDM updates the subscription data of hSMF by Nudm_SDM_Notification (SUPI, Session Management Subscription Data). The hSMF updates the Session Management Subscription Data and acknowledges the UDM by returning an Ack with (SUPI).
1d	The SMF might modify the PDU session. This procedure can also be triggered based on locally configured policy or triggered from the RAN.
1e	RAN indicates to the SMF when the AN resources onto which a QoS Flow is mapped are released irrespective of whether notification control is configured. RAN sends the N2 message to the AMF. The AMF invokes Nsmf_PDUSession_UpdateSMContext procedure. The vSMF sends an Nsmf_PDUSession_Update Request message to the hSMF. The hSMF acknowledges and sends an Nsmf_PDUSession_Update Response message to the vSMF.
2	The SMF reports some subscribed event to the PCF by performing an SMF initiated SM Policy Association Modification procedure.
3	The hSMF invokes the Nsmf_PDUSession_Update Request service operation to the vSMF.
3a	The hSMF initiates an N4 Session Modification procedure with the selected hUPF.
3b	The hUPF acknowledges and sends an N4 Session Modification Response message to the hSMF.
4a	The vSMF sends an Nsmf_PDUSession_UpdateSMContext Response message to the AMF.
4b	The vSMF sends an Nsmf_PDUSession_SMContextStatusNotify message to the AMF.
5	The AMF sends the NAS message containing PDU Session ID and PDU Session Establishment Accept targeted to the UE and the N2 SM information received from the SMF within the N2 PDU Session Request to the RAN.
6	The RAN issues AN specific signalling exchange with the UE that is related with the information received from SMF.
7	The AMF sends the Nsmf_PDUSession_UpdateSMContext Request message to the vSMF. The AMF forwards the N2 SM information received from RAN to the vSMF.
8	The AMF sends the Nsmf_PDUSession_UpdateSMContext Request message to the vSMF. The AMF forwards the N2 SM information received from RAN to the vSMF.
9a	The vSMF initiates an N4 Session Modification procedure with the vUPF. The vSMF provides packet detection, enforcement, and reporting rules to be installed on the vUPF for this PDU session, including AN Tunnel Information, H-CN Tunnel Information and V-CN Tunnel Information.
9b	The vUPF provides an N4 Session Modification Response to the vSMF. After this step, the vUPF delivers any down-link packets to the UE that might have been buffered for this PDU Session.
10	The vSMF sends an Nsmf_PDUSession_UpdateSMContext Response message to the AMF.
11	The UE sends a PDU Session Modification Command Ack message to the RAN.
12	The RAN initiates an N2 NAS Uplink Transfer with the AMF.

Step	Description
13	The vSMF sends an Nsmf_PDUSession_UpdateSMContext Response message to the AMF.
14	The vSMF sends an Nsmf_PDUSession_SMContextStatusNotify message to the AMF.
15	The vSMF responds to the hSMF with an Nsmf_PDUSession_Update response carrying the information like PCO provided by the UE in the SM PDU Session Modification Command Ack message from the UE to the vSMF. The hSMF modifies the PDU session context.
16a	The vSMF initiates an N4 Session Modification procedure with the vUPF. The vSMF provides packet detection, enforcement, and reporting rules to be installed on the vUPF for this PDU session, including AN Tunnel Information, H-CN Tunnel Information and V-CN Tunnel Information.
16b	The vUPF provides an N4 Session Modification Response to the vSMF. After this step, the vUPF delivers any down-link packets to the UE that might have been buffered for this PDU Session.
17	The hSMF initiates an SM Policy Association Modification procedure.

PDU Session Release Procedure

This call flow describes the PDU Session Release procedure for the UPF.

Figure 28: PDU Session Release Call Flow

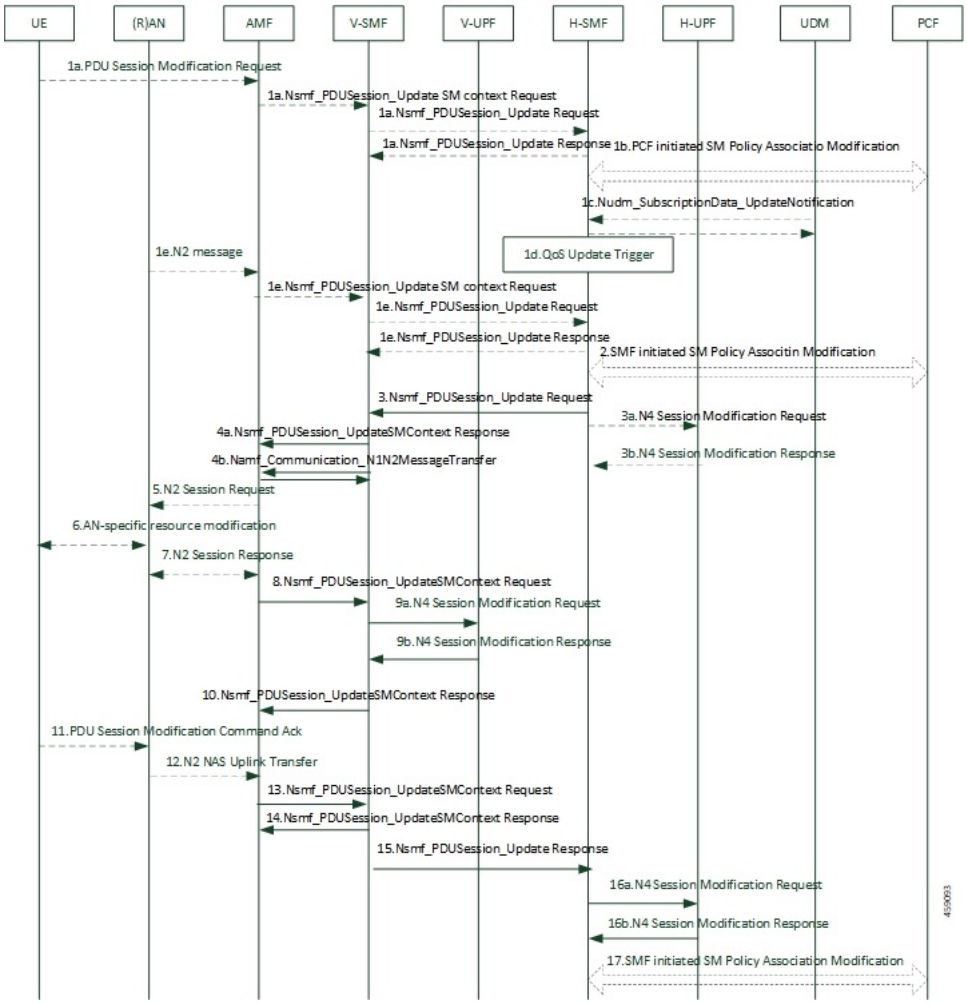


Table 97: PDU Session Release Call Flow Description

Step	Description
1a	• The UE initiates the UE Requested PDU Session Release procedure by the transmission of an NAS message to the AMF. • The AMF invokes the Nsmf_PDUSession_UpdateSMContext service operation and provides the N1 SM container to the SMF together with User Location Information (ULI) received from the RAN. • The vSMF initiates N4 Session Modification to instruct the vUPF to stop forwarding uplink traffic. • The vSMF invokes the Nsmf_PDUSession_Update Request service operation to request the hSMF to release the PDU Session. The hSMF responds to the request immediately.

Step	Description
1b	• The AMF invokes the Nsmf_PDUSession_ReleaseSMContext service operation to request the release of the PDU Session. • The vSMF initiates N4 Session Modification to instruct the vUPF to stop forwarding uplink traffic. • The vSMF initiates the release of the PDU Session at the hSMF by invoking the Nsmf_PDUSession_Release request.
1c	• The PCF invokes an SM Policy Association Termination procedure to request the release of the PDU Session. • The hSMF initiates N4 Session Modification to instruct the hUPF to stop forwarding downlink traffic.
1d	• RAN indicates to the vSMF that the PDU Session-related resource is released when all the QoS Flow(s) of the PDU Session are released. • The vSMF initiates N4 Session Modification to instruct the vUPF to stop forwarding uplink traffic. • The vSMF initiates the Nsmf_PDUSession_Update Request toward the hSMF and the hSMF acknowledges with a response.
1e	The SMF decides to release a PDU session. The hSMF initiates N4 Session Modification to instruct the hUPF to stop forwarding downlink traffic.
1f	The AMF invokes the Nsmf_PDUSession_UpdateSMContext service operation with a release indication to request the release of the PDU session. The vSMF initiates N4 Session Modification to instruct the vUPF to stop forwarding uplink traffic. The vSMF invokes the Nsmf_PDUSession_Update Request toward the hSMF.
2a	The hSMF sends an N4 Session Release Request (N4 Session ID) message to the hUPFs of the PDU session. The hUPFs drop any remaining packets of the PDU session and release all tunnel resource and contexts associated with the N4 Session.
2b	The hUPF(s) acknowledges the N4 Session Release Request by the transmission of an N4 Session Release Response message to the hSMF.
3a	The SMF responds to the AMF with the Nsmf_PDUSession_UpdateSMContext response.
4a	The vSMF sends an N4 Session Release request to the vUPF.
4b	The vUPF acknowledges and sends an N4 Session Release response to the hSMF.
5a	The vSMF sends an Nsmf_PDUSession_ReleaseSMContext Response message to the AMF.
5b	The vSMF sends an Nsmf_PDUSession_UpdateSMContext Response message to the AMF.
5c	The N1N2 Message Transfer procedure occurs between the AMF and vSMF.

Step	Description
6	The hSMF includes the N2 SM Resource Release request in the message sent to the AMF, then the AMF transmits the NAS message to the UE.
7	When the RAN receives an N2 SM request to release the AN resources associated with the PDU session, it issues AN specific signalling exchanges with the UE to release the corresponding AN resources.
8	If the RAN receives an N2 SM request to release the AN resources, it acknowledges the N2 SM Resource Release Request by sending an N2 SM Resource Release Ack message to the AMF.
9	The AMF sends an Nsmf_PDUSession_UpdateSMContext Response message to the vSMF.
10	The vSMF responds to the AMF with an Nsmf_PDUSession_UpdateSMContext response.
11	The UE acknowledges the PDU Session Release Command by sending a NAS message over the RAN.
12	The AMF invokes the Nsmf_PDUSession_UpdateSMContext to the vSMF.
13	The SMF responds to the AMF with an Nsmf_PDUSession_UpdateSMContext response.
14	The vSMF responds to the hSMF with an Nsmf_PDUSession_Update Request invoked at step 3a and confirms the PDU session release.
15	The vSMF initiates an N4 Session Modification procedure with the vUPF. The vSMF provides packet detection, enforcement, and reporting rules to be installed on the vUPF for this PDU session, including AN Tunnel Information, H-CN Tunnel Information and V-CN Tunnel Information.
15a	The hSMF releases the SM policy control association with the PCF by invoking the SM Policy Association Termination procedure.
15b - 15c	In case the PDU Session Release is HPLMN-initiated, the hSMF releases the corresponding User Plane resources.
15d	The hSMF invokes the Nudm_UECM_Deregistration service operation.
16a	The hSMF requests the vSMF to release all contexts associated with the PDU session by invoking the Nsmf_PDUSession_StatusNotify (Release) operation.
16b	The vSMF requests the AMF to release all contexts associated with the PDU Session by invoking the Nsmf_PDUSession_SMContexStatusNotify (Release). The AMF releases the association between the vSMF ID and the PDU Session ID.

5G to 4G Handover

This call flow describes the 5G to 4G handover.

Figure 29: 5G to 4G Handover Call Flow

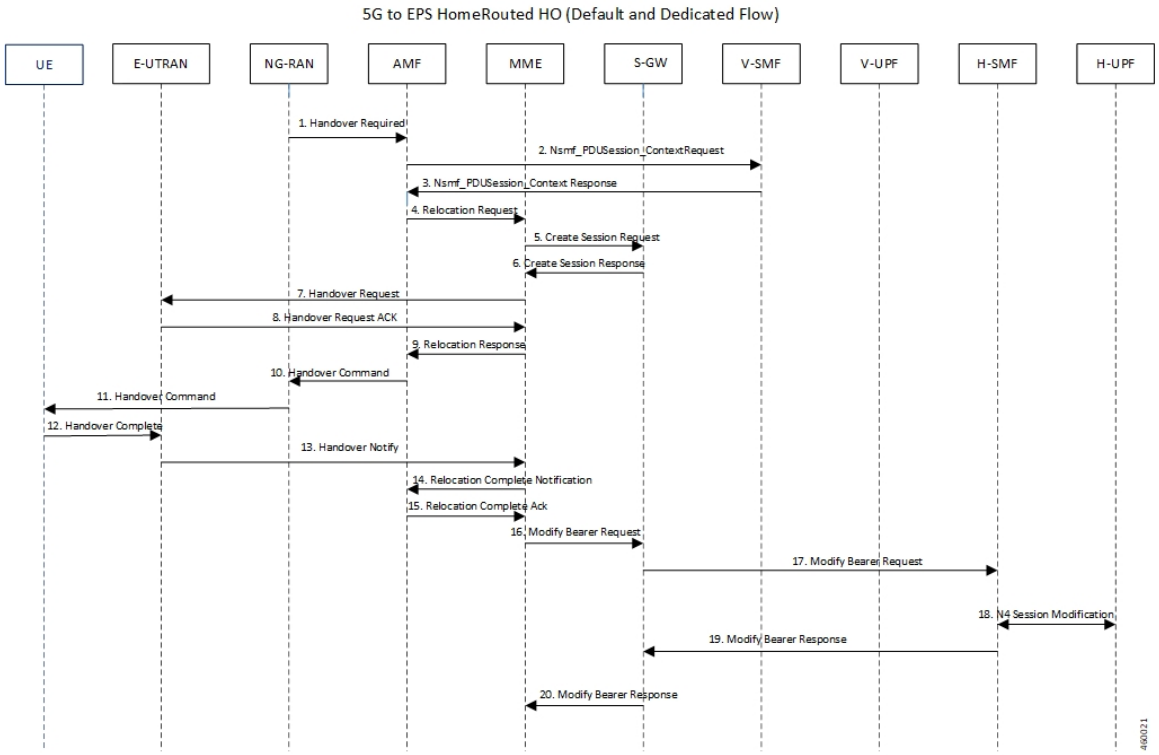


Table 98: 5G to 4G Handover Call Flow Description

Step	Description
1	After the 5G session is established, the NG-RAN initiates the handover process by sending the Handover Required message to the AMF.
2	The AMF invokes the Nsmf_PDUSession_Context Request and sends it to the vSMF.
3	The vSMF sends a Nsmf_PDUSession_Context Response to the AMF.
4	The AMF sends a Relocation Request to the MME.
5	The MME sends Create Session Request to the S-GW.
6	The S-GW sends a Create Session Response message back to the MME.
7	The MME sends a Handover Request message to E-UTRAN.
8	The E-UTRAN acknowledges and sends a Handover Request ACK message back to the MME.
9	The MME sends the Relocation Response message to the AMF.
10	The AMF sends a Handover Command message to the NG-RAN.

Step	Description
11	The NG-RAN commands the UE to handover to the target access network by sending the Handover Command.
12	The UE responds to the E-UTRAN with a Handover Complete message, and the uplink data path is established.
13	The E-UTRAN notifies the MME that the UE is handed over to the NG-RAN.
14	The MME sends a Relocation Complete Notification message to the AMF.
15	The AMF acknowledges and sends a Relocation Complete Ack message to the MME.
16	The MME sends a Modify Bearer Request message to the S-GW.
17	The S-GW forwards the Modify Bearer Request message to the hSMF.
18	The hSMF initiates an N4 Session Modification procedure with the hUPF.
19	The hSMF responds to the S-GW with a Modify Bearer Response message, and the downlink data path is established.
20	The S-GW sends the Modify Bearer Response message to the MME.

4G to 5G Handover

This call flow describes the 4G to 5G handover.

Figure 30: 4G to 5G Handover Call Flow

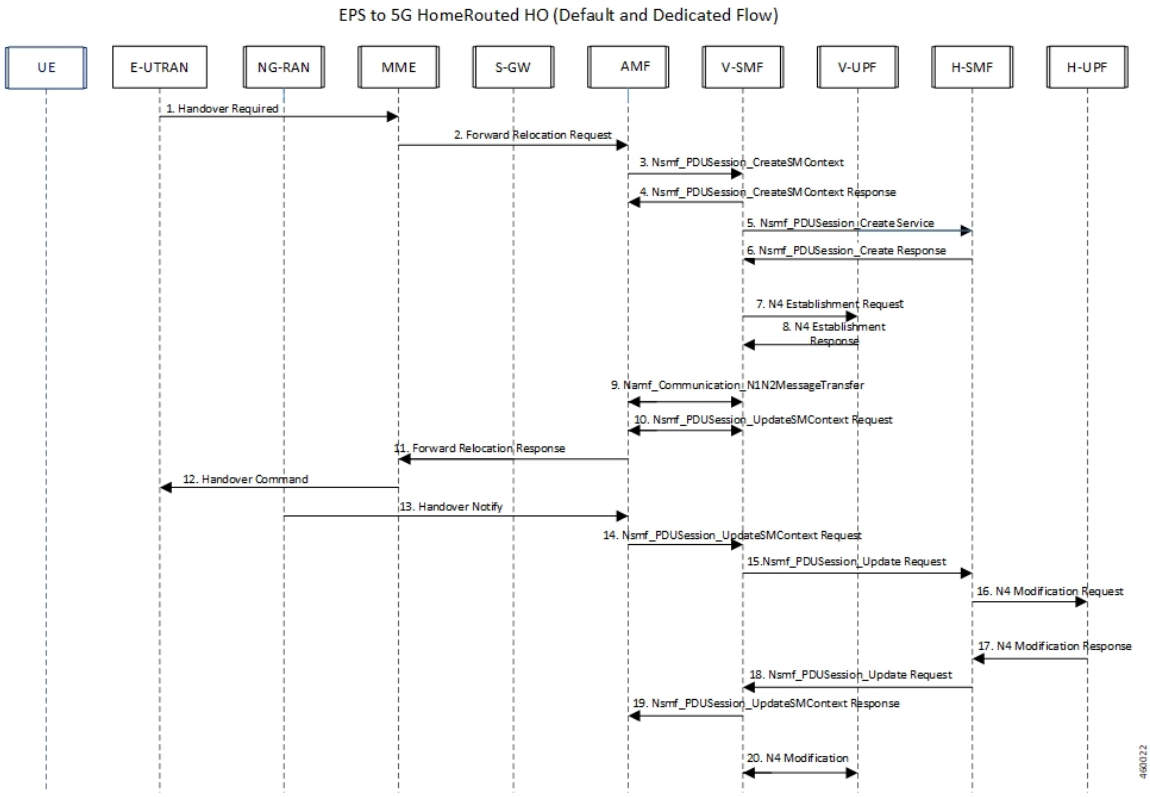


Table 99: 4G to 5G Handover Call Flow Description

Steps	Description
1	After the 4G session is established, the E-UTRAN initiates the handover process by sending the Handover Required message to the MME.
2	The MME sends a Forward Relocation Request to the AMF.
3	The AMF invokes the Nsmf_PDUSession_CreateSMContext Request and sends it to the vSMF.
4	The vSMF sends a Nsmf_PDUSession_CreateSMContext Response to the AMF.
5	The vSMF sends a Nsmf_PDUSession_Create Service message to create a new PDU Session in the hSMF.
6	The hSMF responds with a Nsmf_PDUSession_Create Response message.
7	The vSMF sends an N4 Establishment Request to the vUPF.
8	The vUPF acknowledges by sending an N4 Establishment Response.
9	The vSMF sends an Namf_Communication_N1N2MessageTransfer message to the AMF.
10	The AMF sends the Nsmf_PDUSession_UpdateSMContext Request message to the vSMF.

Steps	Description
11	The AMF sends a Forward Relocation Response message to the MME.
12	The MME sends the Handover Command to the E-UTRAN.
13	The NG-RAN notifies the AMF that the UE is handed over to the NG-RAN.
14	The AMF sends the Nsmf_PDUSession_UpdateSMContext Request message to the vSMF.
15	The vSMF initiates the Nsmf_PDUSession_Update Request toward the hSMF.
16	The hSMF sends an N4 Modification Request with PDRs, FARs, QERs, or URRs to the hUPF.
17	The hUPF responds with an N4 Modification Response message.
18	The hSMF invokes the Nsmf_PDUSession_Update Request service operation to the vSMF.
19	The vSMF sends an Nsmf_PDUSession_UpdateSMContext Response message to the AMF.
20	The vSMF initiates an N4 Modification procedure with the vUPF.



Note To support the roaming status functionality, the Roaming Status IE is sent in Subscriber Params during PDU Session Establishment or during handover.

Inter-PLMN 5G HO from Homer to Roamer

This call flow describes the inter-PLMN 5G handover (HO) procedure with v-SMF insertion.

The UE moves from a homer in H-PLMN to a roamer in V-PLMN. When V-SMF is newly inserted, SMF transitions to H-SMF.

Figure 31: Inter-PLMN 5G HO from Homer to Roamer Call Flow

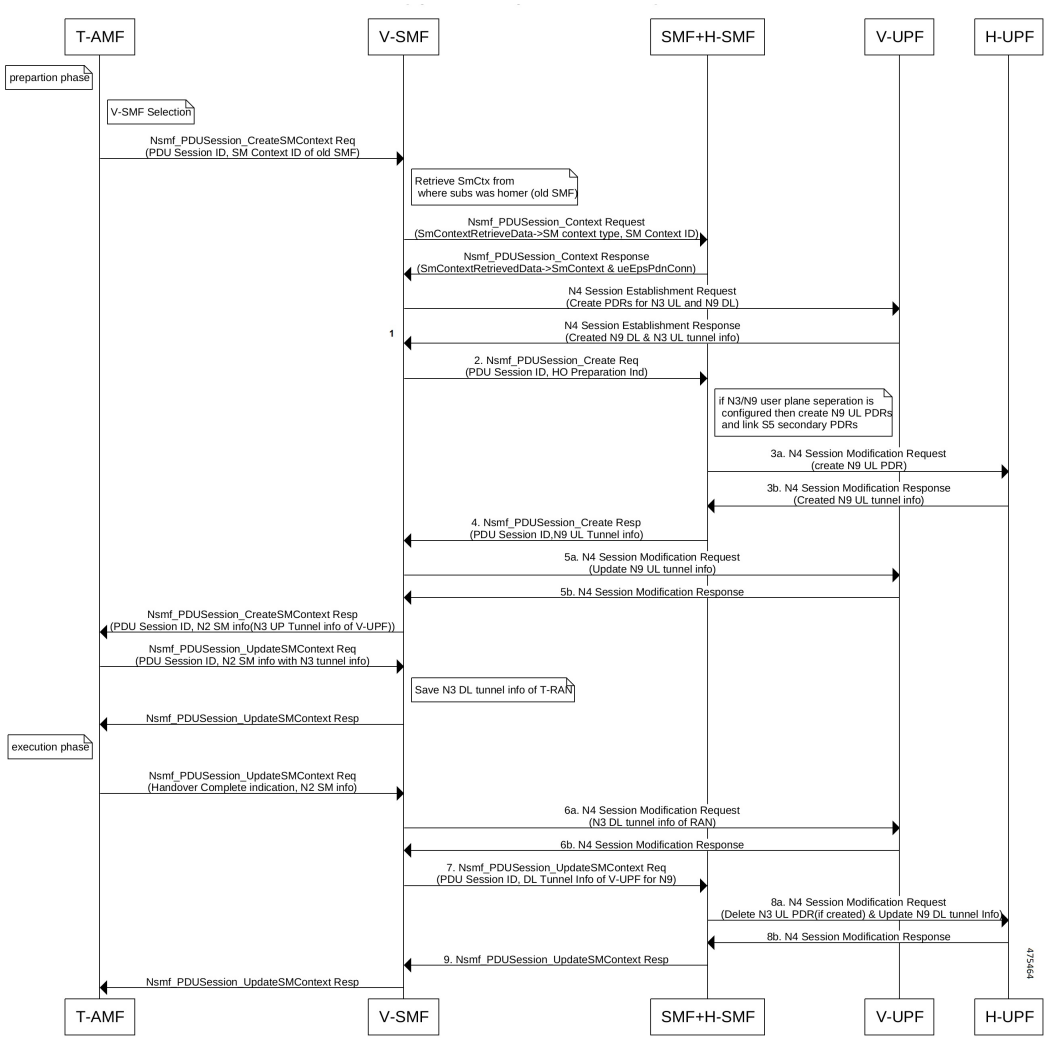


Table 100: Inter-PLMN 5G HO from Homer to Roamer Call Flow Description

Steps	Description
1	UE moves from H-PLMN (5G Homer) to V-PLMN (5G Roamer).
2	V-SMF sends Nsmf_PDUSession_Create Request with HO Preparation Ind info to the H-SMF.
3a	H-SMF sends the following Create PDRs in an N4 Session Modification Request to the H-UPF: • Primary Create PDR with Source Interface Type as N9, Local FTEID having the same FTEID as assigned in S5U Create PDR. • Secondary Create PDR with Source Interface Type as S5-U, Local FTEID having the same FTEID as assigned in S5U Create PDR.
3b	H-UPF sends N4 Session Modification Response with created N9 UL Tunnel info to the H-SMF.

Steps	Description
4	H-SMF sends Nsmf_PDUSession_Create Response with N9 UL Tunnel info to the V-SMF.
5a	V-SMF sends N4 Session Modification Request with N9 UL Tunnel info to the V-UPF.
5b	V-UPF sends N4 Session Modification Response back to the V-SMF.
6a	V-SMF sends N4 Session Modification Request with N3 DL Tunnel Info of RAN to the V-UPF.
6b	V-UPF sends N4 Session Modification Response back to the V-SMF.
7	V-SMF sends Nsmf_PDUSession_UpdateSMContext Request with DL Tunnel info of V-UPF for N9 to the H-SMF.
8a	H-SMF deletes N3 UL PDR (if created), updates N9 DL tunnel Info in an N4 Session Modification Request and sends to the H-UPF. H-SMF sends Create N9 DL FAR, Remove old N3 PDR, Remove old S5 PDR, Remove FAR, Create Core QER, and Update Core PDR in Sx-Modify.
8b	H-UPF sends N4 Session Modification Response back to the H-SMF.
9	H-SMF sends Nsmf_PDUSession_UpdateSMContext Response back to the V-SMF.

Inter-PLMN HO from 5G Roamer to 4G Homer

This call flow describes the inter-PLMN HO from 5G roamer to 4G homer.

Figure 32: Inter-PLMN HO from 5G Roamer to 4G Homer Call Flow

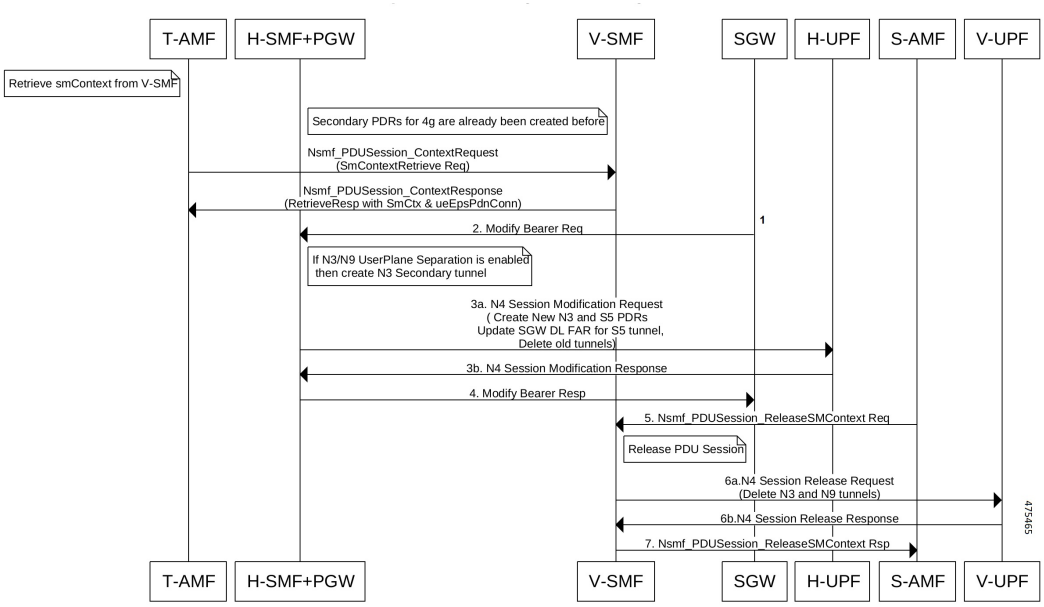


Table 101: Inter-PLMN HO from 5G Roamer to 4G Homer Call Flow Description

Steps	Description
1	UE moves from H-PLMN (5G Roamer) to V-PLMN (4G Homer).

Steps	Description
2	S-GW sends Modify Bearer Request to the H-SMF. Here, if N3/N9 UserPlane Separation is enabled then create N3 secondary tunnel.
3a	H-SMF sends N4 Session Modification Request to H-UPF with the following IEs: Create S5 PDR (old S5-U FTEID), Create N3 PDR (CH bit set for New FTEID), Create S5 DL FAR, Remove old N9 PDR, Remove old S5 PDR, Remove FAR, and Update Core QER.
3b	H-UPF sends N4 Session Modification Response back to the H-SMF.
4	H-SMF sends Modify Bearer Response back to the S-GW.
5	S-AMF sends Nsmf_PDUSession_ReleaseSMContext Request with Release PDU Session to the V-SMF.
6a	V-SMF sends N4 Session Release Request with Delete N3,N9 tunnels to the V-UPF.
6b	V-UPF sends N4 Session Release Response back to the V-SMF.
7	V-SMF sends Nsmf_PDUSession_ReleaseSMContext Response back to the S-AMF.

Standards Compliance

The HR roaming support feature complies with the following standards:

- 3GPP TS 23.502 "Procedures for the 5G System"
- 3GPP TS 29.061 "Interworking between the Public Land Mobile Network (PLMN) supporting packet based services and Packet Data Networks (PDN)"
- 3GPP TS 29.244 16.5.0 "Interface between the Control Plane and the User Plane nodes"

Limitations

The HR roaming support feature has the following limitations:

- Router Advertisement (RA) or Router Solicitation (RS) packets are charged and counted in default QFI PDR on the vUPF.
- No support for QER enforcement policing on the vUPF.
- No support for LI on the vUPF.
- UPF supports two GTPU endpoints at the ingress interface.
- Secondary PDR Optimization support will not be enabled in-between, for an existing call.
- After the upgrade, if the source interface type is not available in Create PDR.Then for existing calls, service selection will not happen based on interface type.
- N9, S5U, and S8U interface types are supported only on the same public IP and not on different IPs.
- Inter-PLMN handover works only with Secondary PDR optimization.

- If any Access-side PDR handling fails, then UPF reports only the failed PDR ID and not both PDRs (primary and secondary).

Configuring the HR Roaming Support for UPF

This section describes how to configure the HR roaming support feature for UPF.

Configure Buffering Support of Visitors Calls on vUPF

To configure the buffering support of visitors calls on the vUPF, use the following CLI commands:

```
config
   user-plane [converged mode | buffered-packet-count [ instance-limit
instance_limit_value { session-limit session_limit_value} | { session-limit
session_limit_value { instance-limit instance_limit_value} ] ]
   exit
```

NOTES:

- **buffering-packet-count**: Configure max session and instance limit for buffering the packets.
- **instance-limit** *instance_limit_value*: Configures maximum number of packets to buffer for all session per SessMgr instance. The default range is 1 to 10000.
- **session-limit** *session_limit_value*: Configures maximum number of packets to buffer per session. The default range is 1 to 255.

Verify the Buffering Support of Visitors Calls on vUPF

To verify the support of buffering limit for traffic on visitor calls on the vUPF, use the **show configuration** CLI command.

The following code is a sample output of the CLI command.

```
[local]qvpc-si# show configuration
...
...
    #exit
  user-plane buffered-packet-count session-limit 5 instance-limit 10
  context ingress
…
…
```

Configure the GTP-U Service and N9 Interface Association

To configure the association of GTP-U Service and N9 Interface, use the following CLI commands:

```
config
   context ingress
     user-plane-service user-plane-service
           associate gtpu-service service_name [ cp-tunnel | pgw-ingress |
 sgw-egress | sgw-ingress | upf-egress | upf-ingress | interface-type [
n9 | s5u | s8u | n3 | n9-s5u-s8u ] ]
           no associate gtpu-service upf-egress
     exit
```

```
   exit
exit
```

NOTES:

- **upf-egress**: Configure the interface type as UPF egress used for N9 interface.
- **interface-type [n9 | s5u | s8u | n3 | n9-s5u-s8u]**: Configure the desired GTP-U ingress interface type.

  Use the **n9-s5u-s8u** interface type a when single public IP address is used for N9, S5u and S8u interface types.
- **no associate gtpu-service upf-egress**: Configure to remove the GTP-U service and N9 interface association.

Verify the GTP-U Service and N9 Interface Association

To verify the association of GTP-U Service and N9 Interface, use the **show user-plane-service all** CLI command.

The following code is a sample output of the CLI command.

```
[local]qvpc-si# show user-plane-service all

Service name                        : user-plane-service
  Service-Id                        : 6
  Context                           : ingress
  Status                            : STARTED
  UPF Ingress GTPU Service          : sx-gtpu-service
  UPF Egress GTPU Service           : sx-upf_egress_gtpu
  SGW Ingress GTPU Service          : sx-sgw_ingress_gtpu
  SGW Egress GTPU Service           : sx-sgw_egress_gtpu
….
[local]qvpc-si# show configuration context ingress
config
  context ingress
....
...
   user-plane-service user-plane-service
      associate gtpu-service sx-gtpu-service upf-ingress
      associate gtpu-service sx-upf_egress_gtpu upf-egress
      associate gtpu-service sx-sgw_ingress_gtpu sgw-ingress
      associate gtpu-service sx-sgw_egress_gtpu sgw-egress
      associate gtpu-service up-gtpu cp-tunnel
      associate sx-service sxu
      associate control-plane-group g1
    exit
...
```

Roaming Status in Sub-Params IE

The following is an example of the Sub-Params IE that displays the roaming status:

```
SUBSCRIBER PARAMS:
Type: 226
Rat-Type: 3
SGSN-Address: 20.20.20.17
ULI: ECGI=216354001234c800
GGSN-Address: 20.20.20.3
Roaming-Status: Homer/Roamer/Visitor-hr/Visitor-LBO
```

Source Interface Type IE

The Source Interface Type IE indicates the 3GPP interface type of the source interface.

The 3GPP Interface Type IE indicates the 3GPP interface type of the source interface within the PDR IE, or the 3GPP interface type of the Destination Interface within the FAR IE.

The following is an example of Create PDR that displays Source Interface Type:

```
CREATE PDR:
    Type: 1
    Value:
PDR ID:
    Type: 56
    Value:
        RULE ID: 0x0002
PRECEDENCE:
    Type: 29
    Value: 0x00000001

PDI:
Type: 2
Value:
SOURCE INTERFACE:
    Type: 20
    Value: CORE (1)
Traffic Endpoint ID:
    Type: 131
    Value: 0x0001
SOURCE INTERFACE TYPE:
    Type: 160
    Value: N9 (0x0F)
FAR ID:
    Type: 108
    Value: 0x8002
RULEBASE:
    Type: 207
    Value:
QER ID:
    Type: 109
    Value: 0x80000001
```

Monitoring and Troubleshooting

This section provides troubleshooting information for this feature.

show user-plane-service statistics all

To monitor the statistics for the user plane service, use the following CLI command:

show user-plane-service statistics all

The following is a sample output:

```
[local]qvpc-si# show user-plane-service statistics all

…
PDNs By PLMN-Type:
 Homer Subscriber PDNs:
    Active:                          0    Setup:                               0
```

```
    Released:                        0
 Roamer Subscriber PDNs:
    Active:                             0    Setup:                                0
    Released:                        0
  Visited-Home-Routed Subscriber PDNs:
    Active:                             0    Setup:                                0
    Released:                        0
  Visited-Local-Breakout Subscriber PDNs:
    Active:                             0    Setup:                                0
    Released:                        0
…
Inter-PLMN Handover Statistics:
   Attempted  : 0                    Succeeded  : 0
   Failed.        : 0

  Intra-PLMN Handover Statistics:
   Attempted   : 0                    Succeeded  : 0
    Failed.        : 0
…
  PDNs Rejected By Reason:
…
…
  Data Statistics Related To Paging:
    Packets Buffered:                   5    Bytes Buffered:                   420
    Packets Discarded:                  1    Bytes Discarded:                   84

  Total Data Statistics:
```

show subscribers user-plane-only full all

To see all the subscribers using the user plane service, use the following CLI command:

show subscribers user-plane-only full all

The following is a sample output:

```
….
Converged Session: No                Converged Peer Callid:      n/a
Visited Call: Yes
Subscriber Parameters:
…..
Interface Type: N4
….
Subscriber Parameters:
IMSI: 123456789012345
….
Old Rat Type: 6
…
Session-ID: 1414146500000001
Roaming Status: Homer
Old Roaming Status: Homer
…
```

show subscribers user-plane-only callid pdr full all

The **show subscribers user-plane-only callid** *call_id* **pdr full all** command displays all PDR information.

The Source Interface Type field displays the Source Interface Type IE received from SMF.

The following is a sample output of this command:

show subscribers user-plane-only callid *call_id* **pdr full all**

```
show subscribers user-plane-only callid 0098e4b0 pdr full all
Callid: 0098e4b0
Interface Type: N4
…
Matched Bytes:            0
Matched Packets:          0
Precedence:               0
Source Interface:         Access
Source Interface Type:    N9
Local F-TEID:    0x16001
Local F-TEID IP Addr:
  IPv4:          20.20.20.15
  IPv6:          N/A
…
Bearer Level Info-id:   N/A
Secondary Pdr Info:
   Pdr ID :   2 / NA
   Local F-TEID:    0x1a001
   Local F-TEID IP Addr:
         IPv4:          20.20.20.15
         IPv6:          N/A
   QoS Flow Identifier: 1 / NA
   Source Interface Type:  S5U
```

show user-plane-service gtpu statistics gtpu-service n9-egress

The following is a sample output of this command for visiting UPF:

```
[local]Test-UPF4# show user-plane-service gtpu statistics gtpu-service n9-egress

Total Data Stats:
Uplink Packets: 10 Uplink Bytes: 2400
Downlink Packets: 10 Downlink Bytes: 2400
Packets Discarded: 0 Bytes Discarded: 0

QoS Stats:

QCI 1:
Uplink Packets: 0 Uplink Bytes: 0
Downlink Packets: 0 Downlink Bytes: 0
Packets Discarded: 0 Bytes Discarded: 0

QCI 2:
Uplink Packets: 8 Uplink Bytes: 1920
Downlink Packets: 8 Downlink Bytes: 1920
Packets Discarded: 0 Bytes Discarded: 0

QCI 3:
Uplink Packets: 0 Uplink Bytes: 0
Downlink Packets: 0 Downlink Bytes: 0
Packets Discarded: 0 Bytes Discarded: 0

QCI 4:
Uplink Packets: 0 Uplink Bytes: 0
Downlink Packets: 0 Downlink Bytes: 0
Packets Discarded: 0 Bytes Discarded: 0

QCI 5:
Uplink Packets: 0 Uplink Bytes: 0
Downlink Packets: 0 Downlink Bytes: 0
Packets Discarded: 0 Bytes Discarded: 0

QCI 6:
```

```
Uplink Packets: 0 Uplink Bytes: 0
Downlink Packets: 0 Downlink Bytes: 0
Packets Discarded: 0 Bytes Discarded: 0

QCI 7:
Uplink Packets: 0 Uplink Bytes: 0
Downlink Packets: 0 Downlink Bytes: 0
Packets Discarded: 0 Bytes Discarded: 0

QCI 8:
Uplink Packets: 0 Uplink Bytes: 0
Downlink Packets: 0 Downlink Bytes: 0
Packets Discarded: 0 Bytes Discarded: 0

QCI 9:
Uplink Packets: 2 Uplink Bytes: 480
Downlink Packets: 2 Downlink Bytes: 480
Packets Discarded: 0 Bytes Discarded: 0

QCI 65:
Uplink Packets: 0 Uplink Bytes: 0
Downlink Packets: 0 Downlink Bytes: 0
Packets Discarded: 0 Bytes Discarded: 0

QCI 66:
Uplink Packets: 0 Uplink Bytes: 0
Downlink Packets: 0 Downlink Bytes: 0
Packets Discarded: 0 Bytes Discarded: 0

QCI 69:
Uplink Packets: 0 Uplink Bytes: 0
Downlink Packets: 0 Downlink Bytes: 0
Packets Discarded: 0 Bytes Discarded: 0

QCI 70:
Uplink Packets: 0 Uplink Bytes: 0
Downlink Packets: 0 Downlink Bytes: 0
Packets Discarded: 0 Bytes Discarded: 0

QCI 80:
Uplink Packets: 0 Uplink Bytes: 0
Downlink Packets: 0 Downlink Bytes: 0
Packets Discarded: 0 Bytes Discarded: 0

QCI 82:
Uplink Packets: 0 Uplink Bytes: 0
Downlink Packets: 0 Downlink Bytes: 0
Packets Discarded: 0 Bytes Discarded: 0

QCI 83:
Uplink Packets: 0 Uplink Bytes: 0
Downlink Packets: 0 Downlink Bytes: 0
Packets Discarded: 0 Bytes Discarded: 0

Non-Std QCI(Non-GBR):
Uplink Packets: 0 Uplink Bytes: 0
Downlink Packets: 0 Downlink Bytes: 0
Packets Discarded: 0 Bytes Discarded: 0
Uplink Packets: 0 Uplink Bytes: 0
Downlink Packets: 0 Downlink Bytes: 0
Packets Discarded: 0 Bytes Discarded: 0

Total uplink packets GBR QCI's: 8
Total uplink Bytes GBR QCI's: 1920
```

```
Total Downlink packets GBR QCI's: 8
Total Downlink Bytes GBR QCI's: 1920
Total uplink packets Non-GBR QCI's: 2
Total uplink Bytes Non-GBR QCI's: 480
Total Downlink packets Non-GBR QCI's: 2
Total Downlink Bytes Non-GBR QCI's: 480

Path Management Messages:
Echo Request Rx: 29 Echo Response Rx: 10
Echo Request Tx: 10 Echo Response Tx: 29
SuppExtnHdr Tx: 0 SuppExtnHdr Rx: 0

Peer Stats:
Total GTPU Peers: 1
Total GTPU Peers with Stats: 1

Tunnel Management Messages:
Error Indication Tx: 0
Error Indication Rx: 0
Error Indication Rx Discarded: 0
```

show user-plane-service gtpu statistics gtpu-service n9-ingress

The following is a sample output of this command for roaming UPF:

```
[local]Test-UPF3# show user-plane-service gtpu statistics gtpu-service n9-ingress

Total Data Stats:
Uplink Packets: 25 Uplink Bytes: 6000
Downlink Packets: 24 Downlink Bytes: 5760
Packets Discarded: 0 Bytes Discarded: 0

QoS Stats:

QCI 1:
Uplink Packets: 0 Uplink Bytes: 0
Downlink Packets: 0 Downlink Bytes: 0
Packets Discarded: 0 Bytes Discarded: 0

QCI 2:
Uplink Packets: 19 Uplink Bytes: 4560
Downlink Packets: 18 Downlink Bytes: 4320
Packets Discarded: 0 Bytes Discarded: 0

QCI 3:
Uplink Packets: 0 Uplink Bytes: 0
Downlink Packets: 0 Downlink Bytes: 0
Packets Discarded: 0 Bytes Discarded: 0

QCI 4:
Uplink Packets: 0 Uplink Bytes: 0
Downlink Packets: 0 Downlink Bytes: 0
Packets Discarded: 0 Bytes Discarded: 0

QCI 5:
Uplink Packets: 0 Uplink Bytes: 0
Downlink Packets: 0 Downlink Bytes: 0
Packets Discarded: 0 Bytes Discarded: 0

QCI 6:
Uplink Packets: 0 Uplink Bytes: 0
Downlink Packets: 0 Downlink Bytes: 0
Packets Discarded: 0 Bytes Discarded: 0
```

```
QCI 7:
Uplink Packets: 0 Uplink Bytes: 0
Downlink Packets: 0 Downlink Bytes: 0
Packets Discarded: 0 Bytes Discarded: 0

QCI 8:
Uplink Packets: 0 Uplink Bytes: 0
Downlink Packets: 0 Downlink Bytes: 0
Packets Discarded: 0 Bytes Discarded: 0

QCI 9:
Uplink Packets: 6 Uplink Bytes: 1440
Downlink Packets: 6 Downlink Bytes: 1440
Packets Discarded: 0 Bytes Discarded: 0

QCI 65:
Uplink Packets: 0 Uplink Bytes: 0
Downlink Packets: 0 Downlink Bytes: 0
Packets Discarded: 0 Bytes Discarded: 0

QCI 66:
Uplink Packets: 0 Uplink Bytes: 0
Downlink Packets: 0 Downlink Bytes: 0
Packets Discarded: 0 Bytes Discarded: 0

QCI 69:
Uplink Packets: 0 Uplink Bytes: 0
Downlink Packets: 0 Downlink Bytes: 0
Packets Discarded: 0 Bytes Discarded: 0

QCI 70:
Uplink Packets: 0 Uplink Bytes: 0
Downlink Packets: 0 Downlink Bytes: 0
Packets Discarded: 0 Bytes Discarded: 0

QCI 80:
Uplink Packets: 0 Uplink Bytes: 0
Downlink Packets: 0 Downlink Bytes: 0
Packets Discarded: 0 Bytes Discarded: 0

QCI 82:
Uplink Packets: 0 Uplink Bytes: 0
Downlink Packets: 0 Downlink Bytes: 0
Packets Discarded: 0 Bytes Discarded: 0

QCI 83:
Uplink Packets: 0 Uplink Bytes: 0
Downlink Packets: 0 Downlink Bytes: 0
Packets Discarded: 0 Bytes Discarded: 0

Non-Std QCI(Non-GBR):
Uplink Packets: 0 Uplink Bytes: 0
Downlink Packets: 0 Downlink Bytes: 0
Packets Discarded: 0 Bytes Discarded: 0
Uplink Packets: 0 Uplink Bytes: 0
Downlink Packets: 0 Downlink Bytes: 0
Packets Discarded: 0 Bytes Discarded: 0

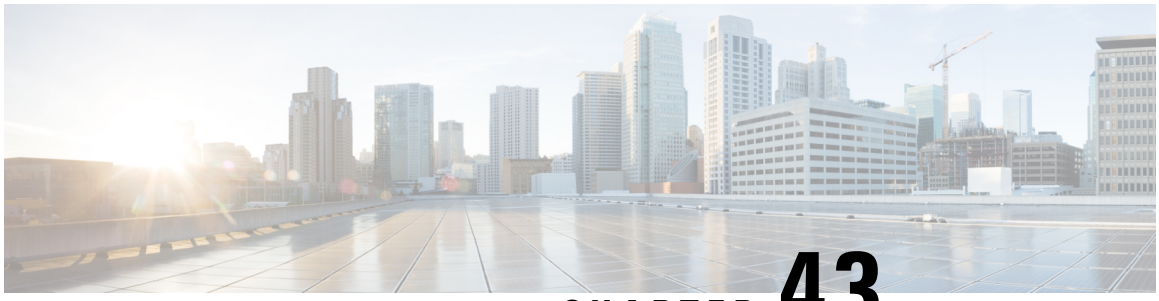
Total uplink packets GBR QCI's: 19
Total uplink Bytes GBR QCI's: 4560
Total Downlink packets GBR QCI's: 18
Total Downlink Bytes GBR QCI's: 4320
Total uplink packets Non-GBR QCI's: 6
```

```
Total uplink Bytes Non-GBR QCI's: 1440
Total Downlink packets Non-GBR QCI's: 6
Total Downlink Bytes Non-GBR QCI's: 1440

Path Management Messages:
Echo Request Rx: 10 Echo Response Rx: 31
Echo Request Tx: 31 Echo Response Tx: 10
SuppExtnHdr Tx: 0 SuppExtnHdr Rx: 0

Peer Stats:
Total GTPU Peers: 1
Total GTPU Peers with Stats: 1

Tunnel Management Messages:
Error Indication Tx: 0
Error Indication Rx: 0
Error Indication Rx Discarded: 0
```



CHAPTER 43

Secure Group Tag-based Access Control

- Feature Summary and Revision History, on page 461
- Feature Description, on page 462
- How it Works, on page 464
- Viewing Fastpath Stream Details, on page 464
- User Access Control through SGACL, on page 465
- How SGACL for ISE Integration Works, on page 467
- Secondary ISE server integration on UPF, on page 471
- SGACL Configuration for ISE Integration, on page 474
- Monitoring and Troubleshooting, on page 478

Feature Summary and Revision History

Summary Data

Table 102: Summary Data

Applicable Products or Functional Area	5G-UPF
Applicable Platforms	VPC-SI
Feature Default Setting	Enabled–Always-on
Related Changes in this Release	Not Applicable

Revision History

Revision Details	Release
First introduced.	2023.02.0

Feature Description

Table 103: Feature History

Feature Name	Release Information	Description
Security Group Tag for Cisco Identity Service Engine (ISE) Integration on UPF	2023.03	To support the additional RADIUS attributes, UPF supports ISE integration for handling the Security Group Tag (SGT) received from SMF. The SMF receives the SGT from the RADIUS server. Then, the SMF sends the SGT over the Sx or N4 interface to UPF during Session Establishment Request. The creation of SGT is according to the static policy on Cisco ISE or SMF and the UPF requires inserting the SGT into the CMD header on uplink packets. **Default Setting**: Not Applicable
User Access Control through Secure Group Tag-based Access Control List (SGACL)	2024.03.0	UPF supports Cisco ISE integration for SGACL enforcement on the downlink packets. UPF fetches the SGACL matrix from ISE through an API query based on the Destination SGT (D-SGT). The D-SGT is received over Sx or N4 interface from SMF. Then, UPF applies the SGACLs based on the the D-SGT and Source SGT (S-SGT) mapping on the downlink packets. Hence, the policy enforcement from Cisco ISE is enabled. Cisco ISE, as a centralized policy engine, provides a unified policy management experience for the other Cisco packet core elements. This feature introduces the new CLI commands **ise-server-profile** *profile_name* to define the ISE server profile and **associate ise-server-profile name***server_profile_name* to associate the ISE server profile within a UPF service. **Default Setting**: Disabled – Configuration Required to Enable

The Security Group Tag (SGT) specifies the privileges of a traffic source within a trusted network. Security Group Access, which is a feature of both Cisco TrustSec and Cisco ISE, automatically generates the SGT when a user adds a security group in TrustSec or ISE.

Note Security Group Tag (SGT) is also referred to as Scalable Group Tag.

The Identity Services Engine (ISE) sends SGT values over the RADIUS interface, that are propagated over the N6 interface. The SGT value needs to be checkpointed for both session recovery and ICSR/GR scenarios for availability to the peer session manager.

Cisco ISE

Cisco Identity Services Engine (ISE) is a centralized policy management platform that provides secure network access to end users and devices. The Cisco ISE server enables the creation and enforcement of security and access policies for endpoint devices that are connected to the routers and switches of an enterprise.

Having Cisco ISE as a centralized policy engine helps in unified policy management experience for the other Cisco packet core elements.

- VLAN segmentation—VLAN segmentation is beneficial in the following aspects:
  - Security—Security Groups that have sensitive data are separated from the rest of the network, decreasing the chances of confidential information breaches.
  - Cost reduction—Cost savings result from the reduced need for expensive network upgrades and more efficient use of the existing network bandwidth.
  - Improved network performance—Dividing Layer 2 networks into multiple logical workgroups reduces unnecessary traffic on the network and boosts performance. VLANs enable the implementation of access and security policies according to specific groupings of users.

Security Group Tag

The Security Group Tag (SGT), also referred to as the Scalable Group Tag, specifies the privileges of a traffic source within a trusted network. Security Group Access, which is a feature of both Cisco TrustSec and Cisco ISE, automatically generates SGT when you add a security group in TrustSec or ISE.

To support SGT for ISE integration:

- SMF receives the SGT from the ISE server.
- SMF updates the SGT towards the UPF using the N4 extensions.
- UPF identifies the device packets and applies the corresponding SGT to the N6 packets.

The SGT has the following values as the Security Group Tag-based Access Control List (SGACL) enforcement is done for the downlink traffic:

- Destination SGT (D-SGT)—The SGT value that UPF receives per subscriber session over the Sx or N4 interface from SMF in the Session Establishment Request. The SMF fetches this information from ISE server over the RADIUS interface.
- Source SGT (S-SGT)—The SGT value that is received in the CMD header of a downlink packet. A wireless LAN controller (WLC) or an access switch inserts this value.

The Identity Services Engine (ISE) sends D-SGT values over the RADIUS interface that are propagated over the N6 interface. The D-SGT value needs to be checkpointed for session recovery and ICSR based recovery.

When the user access control through the SGACL feature is enabled and a subscriber session receives a D-SGT, the SGACL is applied to the downlink packets. The SGACL is applied based on the D-SGT and S-SGT mapping.



Note In this case, the APN ACL isn't applied to these sessions.

Security Group Tag-based Access Control List

Security Group Tag-based Access Control List (SGACL) is an Access Control List (ACL) that controls and manages the authorization of the security group members. SGACLs create SGACL policies, which are represented through an SGT matrix.

Security Group Tag Matrix

SGACLs create SGACL policies. Security Group Tag matrix (SGT matrix), also referred to as permissions matrix, represent these policies in the TrustSec domain. This SGT matrix includes the security group numbers and destination security group numbers. This matrix describes how two endpoints communicate. The applicable policies are Permit and Deny. The contents of SGT matrix and SGACLs are downloaded from ISE server using the REST API.

UPF inserts the D-SGT value for the outgoing uplink packets sent over the N6 interface. UPF receives the S-SGT value for the downlink packets over the N6 interface. This S-SGT value is used for the matrix lookup and is removed while sending the outgoing downlink packets over the N3 interface.

Based on the mapping between the Destination SGT (D-SGT) and Source SGT (S-SGT), the policies are enforced at UPF and an appropriate SGACL is enforced on the downlink packets.

Fastpath Stream

The SGT value that is received in the Session Establishment Request is carried with Outer Header Removal (OHR) option for the uplink stream. This operation enables VPP to add SGT in the corresponding uplink packets towards the N6 interface. The per stream SGT is provisioned in fastpath as part of the **remove-outer-header** operation.

How it Works

Viewing Fastpath Stream Details

You can view the fastpath stream details using the **fastpath stream list allinfo** CLI command.

The following is a sample output of this command:

```
vpp# fastpath stream list allinfo

p ~odc qta teid immediate=0xd, 2817:r104354:d14069, 2831:r51892:d51892) Policer,
(Next: V6:254, V4254), (2829:r753755:d753755, 2828:r753755:d753755) Policer,
(Next: V6: 24, V4 25), (2829:r753758:d753758, 2828:r753758:d753758) hcfUL,
(Next: V6: 22, V4 23), (2827:r9964:d9964, 2830:r189893:d189893,
2831:r51892:d51892) remOuterHdr, (Next: V6: 7, V4 7),
(2817:r104354:d14069, immediate=0x2d000120) L1:Sx35b93
holding queue: Empty  x3-index: 0 dep-id: 0 dup-x3-index: 0 dup-dep-id: 0
```

User Access Control through SGACL

Table 104: Feature History

Feature Name	Release Information	Description
User Access Control through Secure Group Tag-based Access Control List (SGACL)	2024.03.0	UPF supports Cisco ISE integration for SGACL enforcement on the downlink packets. SGACL is an Access Control List (ACL) that controls and manages the authorization of the security group members. UPF fetches the SGACL matrix from ISE through an API query based on the Destination SGT (D-SGT). The D-SGT is received over the Sx or N4 interface from SMF. Then, UPF applies the SGACLs based on the D-SGT and Source SGT (S-SGT) mapping on the downlink packets. Hence, the policy enforcement from Cisco ISE is enabled. This feature introduces the new CLI commands **ise-server-profile** *profile_name* to define the ISE server profile and **associate ise-server-profile name** *server_profile_name* to associate the ISE server profile within a UPF service. **Default Setting**: Disabled – Configuration Required to Enable

Feature Name	Release Information	Description
Security Group Tag (SGT) for Cisco Identity Service Engine (ISE) Integration on UPF	2023.03	To support the additional RADIUS attributes, UPF supports ISE integration for handling the SGT received from SMF. The SMF receives the SGT from the RADIUS server. Then, the SMF sends the SGT over the Sx or N4 interface to UPF during Session Establishment Request. The creation of SGT is according to the static policy on Cisco ISE or SMF and the UPF requires inserting the SGT into the Cisco Meta Data (CMD) header on uplink packets. **Default Setting**: Not Applicable

The Security Group Tag (SGT), also referred to as the Scalable Group Tag, specifies the privileges of a traffic source within a trusted network. Security Group Access automatically generates SGT when you add a security group in TrustSec or ISE. Cisco ISE, as a centralized policy engine, provides a unified policy management experience for the other Cisco packet core elements.

The S-SGT is a 16-bit value that is transmitted in the Cisco Meta Data (CMD) field of a Layer 2 Ethernet Frame. The CMD header is inserted after the ".1Q" tag, if available. If the ".1Q" tag is unavailable, the CMD immediately follows the MAC Source Address.

To support SGT for ISE integration:

- SMF receives the D-SGT from the ISE server.
- SMF updates the D-SGT towards the UPF using the N4 extensions.
- UPF identifies the packets and applies the D-SGT and S-SGT combination to N6 DL packets.

Security Group Tag-based Access Control List (SGACL) is an Access Control List (ACL) that controls and manages the authorization of the security group members. SGACLs create SGACL policies, which are represented through a Security Group Tag matrix (SGT matrix).

The SGT matrix, also referred to as the permissions matrix, represents the SGACL policies in the TrustSec domain. This matrix comprises the security group numbers and destination security group numbers, and describes how the two endpoints communicate. The applicable policies are Permit and Deny. The contents of an SGT matrix and the SGACLs are downloaded from the ISE server using the REST API.

UPF inserts the D-SGT value for the outgoing uplink packets sent over the N6 interface. UPF receives the S-SGT value for the downlink packets over the N6 interface. This S-SGT value is used for the matrix lookup and is removed while sending the outgoing downlink packets over the N3 interface.

Based on the mapping between the Destination SGT (D-SGT) and Source SGT (S-SGT), the policies are enforced at UPF and an appropriate SGACL is enforced on the downlink packets.

UPF supports ISE integration for SGACL enforcement for the downlink packets through the following SGT values:

- Destination SGT (D-SGT)—UPF receives this value per subscriber session over the N4 interface from SMF in the Session Establishment Request. SMF receives the D-SGT from ISE in the RADIUS Access Accept message. For this feature, the SMF must send the SGT to UPF in a proprietary IE on the N4 interface.

- Source SGT (S-SGT)—Is received in the CMD header of a downlink packet. A wireless LAN controller (WLC) or an access switch inserts this value.

The Identity Services Engine (ISE) sends D-SGT values over the RADIUS interface that are propagated over the N6 interface. The D-SGT value needs to be checkpointed for session recovery and ICSR-based recovery.

With the ISE integration for SGACL enforcement feature, the same policy is applied on downlink packets for subscribers with the use of ISE. The communication policies among various source and destination endpoints are managed through Cisco ISE and enforced through UPF.

The policies are enforced at UPF through the Secure Group Tag (SGT) value. UPF inserts the D-SGT value for the outgoing uplink packets. UPF receives the S-SGT value for the downlink packets. This value is used for the matrix lookup and is removed for the outgoing downlink packets. Based on the mapping between S-SGT and D-SGT for a specific downlink packet, an appropriate SGACL is enforced.



Note

- SMF receives the D-SGT from ISE in the Access Accept message. For this feature, the SMF must send the SGT to UPF in a proprietary IE on the N4 interface.
- When you enable the user access control through SGACL and a subscriber session receives a D-SGT, then the SGACL is applied to the downlink packets. In this case, the APN ACL isn't applicable to these subscriber sessions.
- When you don't enable the user access control through SGACL or the UPF doesn't receive the D-SGT during the N4 Session Establishment Request, then the APN ACL is applicable as per the existing configuration.

You can define the ISE server profile through the **ise-server-profile** *profile_name* CLI command and associate the ISE server profile within a UPF service through the **associate ise-server-profile name** *server_profile_name* CLI command.

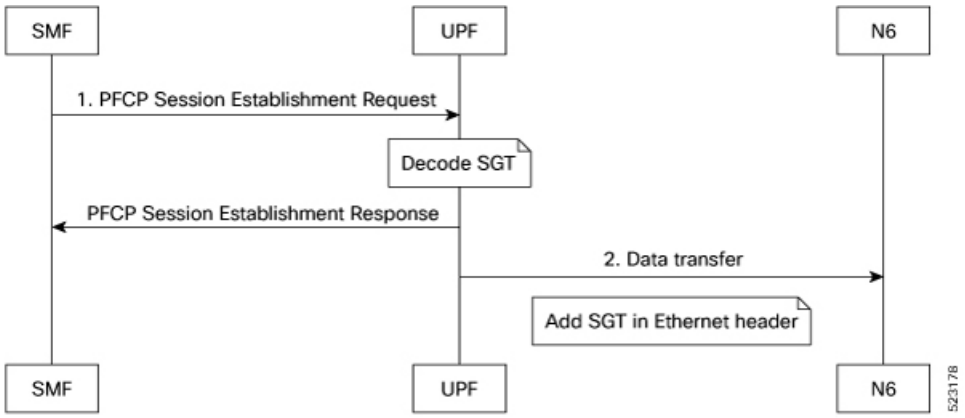
How SGACL for ISE Integration Works

This section describes how SGACL for ISE integration works.

- SMF fetches the corresponding SGT value from ISE over RADIUS and sent on N4 interface to the UPF.
- Each UPF is registered as a Network Access Device (NAD) in ISE server.
- UPF downloads the SGT Matrix and corresponding SGACLs from ISE server using the REST API. UPF applies the SGACL for flow based on the D-SGT to S-SGT mapping.
- A non-real time update is available through periodic or trigger-based pull from UPF.
- The D-SGT value needs to be checkpointed for session recovery and ICSR-based recovery.

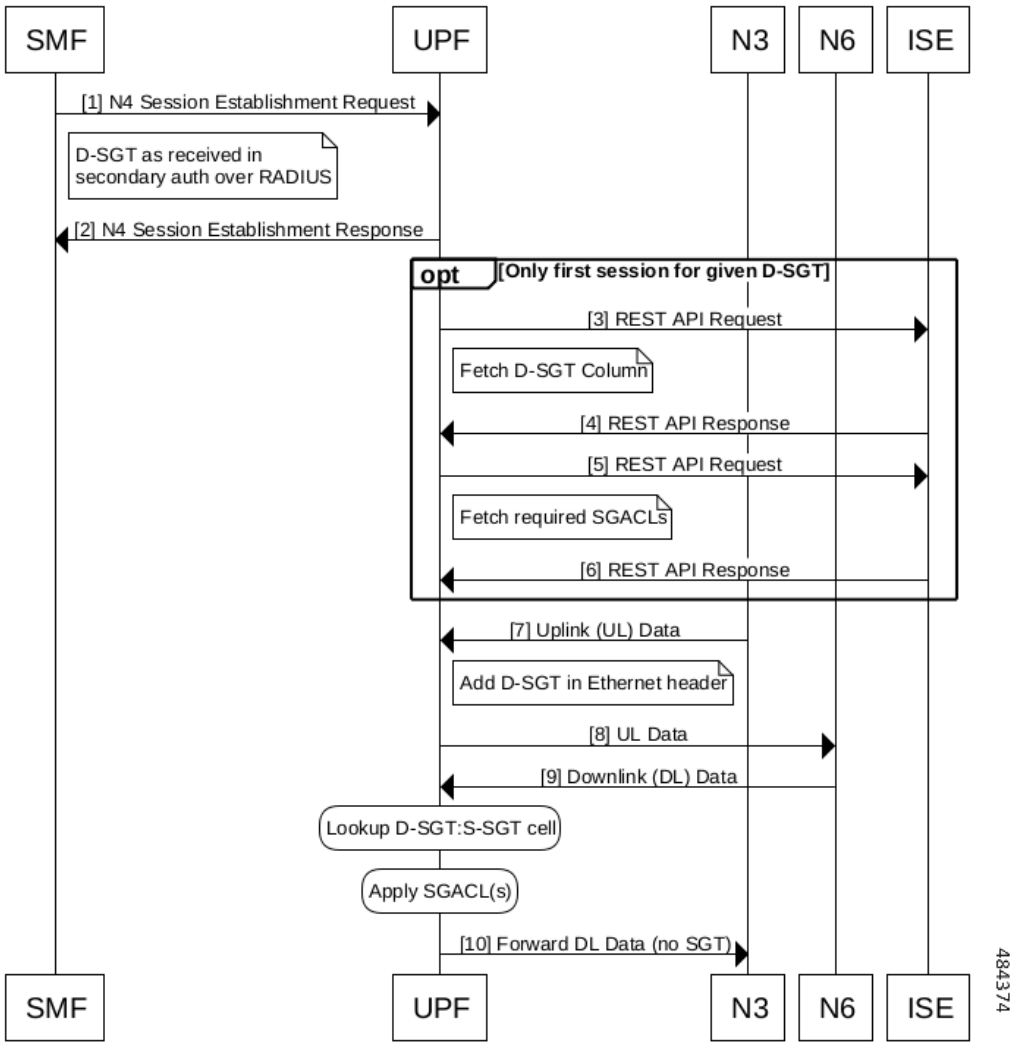
Call Flow

The following figure illustrates the SGT fetch and SGACL enforcement call flow.



Step	Description
1	SMF sends the SGT value in PFCP Session Establishment Request to UPF.
2	UPF decodes the SGT value and adds this value as part of the Cisco Meta Data (CMD) in Ethernet header for all packets sent over the N6 interface.

Figure 33: SGT Fetch and SGACL Enforcement Call Flow



Step	Description
1	SMF sends N4 Session Establishment Request to UPF. As part of this request, the UPF receives the D-SGT value from SMF.
2	UPF sends N4 Session Establishment Response to SMF.
3	For the first session of a specific D-SGT, UPF sends the REST API Request to the ISE server to fetch the D-SGT column values.
4	ISE sends the REST API Response to UPF. This response includes the mapped S-SGT values for the D-SGT and the SGACL matrix information containing the SGACL names along with the refresh timer. The SGT matrix cell entries and SGACLs are stored on UPF. Each D-SGT column has a refresh timer that is configured in ISE.

Step	Description
5	UPF sends the REST API Request to the ISE server to fetch the required SGACLs. ISE sends the REST API Response to UPF with the SGACL definitions.
6	N3 sends the uplink (UL) data to UPF.
7	UPF adds D-SGT in Ethernet header and sends the UL data to the N6 interface.
8	N6 sends the downlink (DL) data to the N3 interface.
9	The UPF checks the D-SGT and S-SGT mapping and applies the SGACLs.
10	UPF forwards the DL data, with no SGT, to the N3 interface.

Limitations

This feature has the following known limitations:

- The PFCP Session Establishment Request must send the D-SGT value over the N4 interface. The D-SGT cannot be changed, but S-SGT can change.
- UPF does not support SGT traffic in the downlink direction.
- As the ACLs can be stacked for a specific SGT-Pair combination, the SGACLs are applied as per the order of ACLs received for a specific SGT-Pair.
- The maximum number of recommended matrix combinations is 150 D-SGT Columns * 150 S-SGT rows, with 255 distinct SGACLs, including the default SGACL.
- The maximum number of SGACLs in a single cell is limited to 16.
- UPF can store only one default SGACL at a time. Hence, when the Default SGACL changes, only the updated SGACLs information is stored on UPF.

Secondary ISE server integration on UPF

Table 105: Feature History

Feature name	Release information	Description
Secondary ISE server support on UPF	2025.03.0	

Feature name	Release information	Description
		This feature allows the network operator to integrate a secondary ISE server on UPF. The secondary ISE server integration provides a failure mechanism to ensure seamless enforcement of SGACLs when the primary ISE server goes down. **Commands enhanced:** • **[secondary-server { ipv4_address** *ipv4_address* **| ipv6-address** *ipv6_address* **}]** : This CLI configures the secondary server IP address. • **[secondary username** *user_name* **password** *password* **]** : This CLI configures the secondary username and password. • **[secondary-certificate** *certificate_path* **]**: This CLI configures the secondary certificate path. • **[secondary-ca-certificate** *ca_certificate_path* **]**: This CLI configures the secondary CA certificate. • **[secondary-key** *key_path* **]**: This CLI configures the secondary key path. **Commands introduced:** • **max-retransmissions** *max_retrans_count* : This CLI configures the maximum number of retries. • **retransmission-timeout** *timeout_duration*: This CLI configures the retransmission timeout. • **backoff-period** *backoff_duration*: This CLI configures the Backoff period.

Feature name	Release information	Description
		Default Settings: Disabled—Configuration to Enable

The ISE server is responsible for enforcing policies on UPF. However, at times, the UPF fails to update or refresh the policy due to a service failure on the ISE server, a failure in reaching the ISE server, or ISE server being down.

Integrating a secondary ISE server allows the UPF to maintain the operational state of the policies and ensure seamless policy enforcement.

How secondary ISE server integration works

Workflow

These stages outline the process of secondary ISE server support on UPF:

1. UPF marks the primary ISE server as "Active". The active primary ISE server receives all the REST API requests.
2. When the UPF tries to connect with the primary ISE server, however, when the primary ISE server fails to respond, the UPF tries to reconnect as per the configured number of retransmissions and the retransmission timeout.
3. UPF marks the current "Active" server as "Inactive", if the reconnect attempts fail. UPF then marks the secondary ISE server as "Active" as per the configuration. To configure the secondary ISE server, see the Configure ISE Server profile section.

   The network operator can select the active ISE server manually using a CLI.

   For more details on the manual configuration, see the Manual switching to active ISE server section.
4. The currently marked "Active" starts receiving the same REST API requests.
5. UPF applies the default policy or policy unavailable treatment, when it is yet to receive the policies from the ISE server.
6. UPF ensures that the secondary ISE server is not marked "Inactive" recently, before marking a server as "Active". It prevents UPF from a looping issue when both servers are unreachable.

   After detecting failure at the Primary ISE server, if the UPF also detects failure at the Secondary ISE server, or vice versa, the UPF should not retry again to the alternative server. This prevents UPF from trying to reach the ISE servers in a loop.

   UPF controls this behavior using the configured Backoff period.

Backoff period

While switching from the primary to the secondary ISE server, or vice versa, the UPF gets into a loop, when both the servers are unreachable. To avoid this, the Backoff period is used as a waiting period before UPF tries to connect with the primary or the secondary ISE server.

The Backoff period prevents excessive network traffic and resource usage caused by continuous retry attempts.

The backoff-period is configurable under the ISE server profile.

Limitations of secondary ISE server integration

These are the known limitations of secondary ISE server integration on UPF:

- ISE server failure is detected only on the next REST API call.
- There is no active monitoring for the health of the primary or secondary ISE server.
- For REST APIs during failover, added latency includes failure detection time (based on retransmission attempts) and connection or reply time from the active ISE server.
- The authentication failures are not expected to trigger a switchover to the secondary ISE server.
- The database synchronization for SGT or SGACLs between primary and secondary ISE servers is assumed, as a switchover can occur between related REST APIs.
- UPF is not responsible for detecting any database version mismatch between the peer servers.
- The periodic or CLI-triggered refresh happens with the current active server, assuming that it is in-sync with the earlier fetched policies.
- UPF will not re-learn existing SGACLs or policies, once the switchover to the active ISE server happens.

  The re-learning adds latency and thus impacts the throughput, since all the flows need to be on-boarded.

SGACL Configuration for ISE Integration

This section describes the procedures to configure SGACL for ISE integration.

Enable API Manager

API manager is a facility that enables API request and response integration through REST APIs.



Note Enabling the API manager is a prerequisite for SGACL integration.

Procedure

Step 1 Log in to the configuration mode.

Step 2 Enter the **require apimgr** command.

Example:

```
config
   require apimgr
end
```

This CLI command is part of the boot configuration to spawn the new procedure.

What to do next

1. Configure ISE Server profile
2. Associate ISE Server Profile

Configure ISE Server profile

These steps help you configure the ISE server profile.

Before you begin

Before configuring the ISE server profile, you must enable the API manager. To know more about configuring the API manager, see the Enable API manager section.

Procedure

Step 1 Use the CLI **[no] ise-server-profile** *profile_name* to create an instance of the ISE server profile under the Context EPC mode.

Example:

```
[local]UPF1(config)# context EPC
[local]UPF1(config-ctx)# ise-server-profile ise_1
[local]UPF1(config-ctx)#
```

Step 2 Use the CLI **bind { ipv4-address** *ipv4_address* **| ipv6-address** *ipv6_address***}** to specify the bind IPv4 or IPv6 address.

Example:

```
[local]UPF1(config-ctx)# bind ipv4-address 192.0.2.1
[local]UPF1(config-ctx)#
```

Step 3 Use the CLI **server { ipv4-address** *ipv4_address* **| ipv6-address** *ipv6_address* **} [secondary-server { ipv4-address** *ipv4_address* **| ipv6-address** *ipv6_address* **}]** to configure primary ISE server and secondary ISE server.

Example:

```
[local]UPF1(config-ctx)# server ipv4-address 192.0.2.254 secondary-server ipv4-address
 198.51.100.1
```

Step 4 Use the CLI **username** *user_name* **encrypted | password** *password* **[secondary username** *user_name* **password** *password* **]** to configure the username and password for the primary ISE server and secondary ISE server.

Example:

```
[local]UPF1(config-ctx)# username ise1 password cisco1
cisco2secondary username ise2 password  cisco2
[local]UPF1(config-ctx)#
```

Note
In the show configuration output, the password will be displayed in the encrypted format.

Step 5 Configure the certificate, CA certificate, and key for the primary and secondary ISE servers.

a) Use the CLIs **certificate** *certificate_path* **[secondary-certificate** *certificate_path* **]** to configure the certificate.

Example:

```
[local]UPF1(config-ctx)# certificate /root/certificate/client/client.cert.pem
secondary-certificate /root/certificate/client/client2.cert.pem
[local]UPF1(config-ctx)#
```

b) Use the CLIs **ca-certificate** *ca_certificate_path* and **[secondary-ca-certificate** *ca_certificate_path* **]** to configure the CA certificate.

Example:

```
[local]UPF1(config-ctx)# ca-certificate /root/certificate/ca.cert.pem
secondary-ca-certificate /root/certificate/ca.cert2.pem
[local]UPF1(config-ctx)#
```

c) Use the CLIs **key** *key_path* **[secondary-key** *key_path* **]** to configure keys.

Example:

```
[local]UPF1(config-ctx)# key /root/certificate/client/client.key.pem
secondary-key /root/certificate/client/client2.key.pem
[local]UPF1(config-ctx)#
```

Step 6 Use the CLI **policy-unavailable-treatment [drop | pass]** to specify the traffic treatment when the SGACL matrix is unavailable for a particular D-SGT.

Example:

```
[local]UPF1(config-ctx)# policy-unavailable-treatment drop
[local]UPF1(config-ctx)#
```

Step 7 Configure maximum retransmissions, retransmission timeout, and backoff period.

a) Use the CLI **max-retransmissions** *max_retrans_count* to configure maximum retransmissions when a server is unresponsive.

Example:

```
[local]UPF1(config-ctx)# max-retransmissions 6
[local]UPF1(config-ctx)#
```

The CLI **max-retransmissions** *max_retrans_count* configures the maximum number of retries from 0 to 15. The default value is 3.

b) Use the CLIs **retransmission-timeout** *timeout_duration* to configure retransmission timeout.

Example:

```
[local]UPF1(config-ctx)# retransmission-timeout 4000
[local]UPF1(config-ctx)#
```

The **retransmission-timeout** *timeout_duration* configures the timeout duration before retransmitting the request to ISE server from 1 to 300 seconds. The default value is 3 seconds.

c) Use the CLI **backoff-period** *backoff_duration* to configure the Backoff-period. Save and exit from the current configuration mode.

Example:

```
[local]UPF1(config-ctx)# backoff-period 400
[local]UPF1(config-ctx)# exit
```

The CLI **backoff-period** *backoff_duration* configures the duration during which, the retries to the ISE server are held back. The default value is 200 seconds.

What to do next

After configuring the ISE server profile, you must associate it with an existing user plane service. For more details, see the Associate ISE server profile section.

Manually switch to active ISE server

These steps allow you to manually switch to the currently active ISE server:

Before you begin

Before manually switching to the currently active ISE server, configure the Secondary ISE server under the ISE server profile. See Configure ISE Server profile section.

Procedure

Execute the CLI **select ise-server { primary | secondary}** in the Execution mode to manually switch to the active ISE server.

Example:

```
[local]laas-si-setup# select ise-server secondary
```

- Post manual switching, if the UPF receives any outstanding REST API responses from the old active ISE server, UPF processes them.
- All the pending retries go to the ISE server, which received the retry requests.
- The command **select** is a hidden keyword.

Associate ISE Server Profile

Once defined, associate the ISE server profile with an existing UPF service configuration.

Before you begin

Configure ISE server profile

Procedure

Step 1 Enter the User Plane Service configuration mode.

Step 2 Associate the defined ISE server profile with an existing UPF service configuration.

Example:

```
user-plane-service UPlane1
      associate ise-server-profile name ise_1
  #exit
```

Refresh D-SGT Column

Each D-SGT column has a refresh timer that is configured in ISE. Based on the refresh timer configuration or through the UPF CLI trigger, the D-SGT column value is fetched from the ISE server again through the REST API query. Based on the API response from ISE, if the version of SGT matrix or SGACL is changed, UPF updates the respective matrix cell or SGACL information locally. The corresponding SGACLs are downloaded after the refresh, as required.

Before you begin

1. Configure ISE Server profile
2. Associate ISE Server profile

Procedure

Step 1 Enter the Exec Mode.

Step 2 Enter the **refresh-sgt-column** *d_sgt* command to trigger the refresh of a D-SGT column by fetching the column values from the ISE server.

Example:

```
refresh-sgt-column d-sgt
```

Note
Although the CLI returns immediately, the policy download in the background takes some time and hence the refresh completion may also take some time.

The following is an example output of the **refresh-sgt-column** *d_sgt* CLI command where the *d_sgt* value is configured as *65535*.

```
refresh-sgt-column 65535
```

Monitoring and Troubleshooting

Verify SGACL with SGT Integration

This section provides information about show commands and their outputs for the SGACL with SGT Integration feature.

show subscriber user-plane-only full all

The output of this command is enhanced to include the following field:

- **SGT Value**: Displays the SGT value.

Fastpath Stream

The SGT value received in Establishment Request is carried with OHR (Outer Header Removal) option for the Uplink stream. This operation enables VPP to add SGT in the corresponding Uplink packets towards N6. The per stream SGT is provisioned in fastpath as part of the remove-outer-header operation.

The following is a sample output of the **fastpath stream** command:

```
vpp# fastpath stream list allinfo

p ~odc qta teid immediate=0xd, 2817:r104354:d14069, 2831:r51892:d51892) Policer,
(Next: V6:254, V4254), (2829:r753755:d753755, 2828:r753755:d753755) Policer,
(Next: V6: 24, V4 25), (2829:r753758:d753758, 2828:r753758:d753758) hcfUL,
(Next: V6: 22, V4 23), (2827:r9964:d9964, 2830:r189893:d189893,
2831:r51892:d51892) remOuterHdr, (Next: V6: 7, V4 7),
(2817:r104354:d14069, immediate=0x2d000120) L1:Sx35b93
holding queue: Empty  x3-index: 0 dep-id: 0 dup-x3-index: 0 dup-dep-id: 0
```

show subscribers user-plane-only full callid *callid_value*

The output of this CLI command is enhanced with the **SGT Value** field for displaying information related to D-SGT for ISE integration on UPF and **SGACL match stats** field for displaying information related to User Access Control through SGACL.

```
show subscribers user-plane-only full callid 00004e3a
   Local SEID      : [0x0004000000000003] 1125899906842627
  Remote SEID     : [0x00000436b7616206] 4633051357702
  State           : Connected

....

  input pkts: 20                                  output pkts: 16
  input bytes: 9246                               output bytes: 11248
  input bytes dropped: 0                          output bytes dropped: 5624
  input pkts dropped: 0                           output pkts dropped: 8

....

  SGT Value: 0x001a

....

QoS-Group Statistics:
QGR Name             Pkts-Down  Bytes-Down Pkts-Up    Bytes-Up   Hits       Match-Bypassed
  FP-Down(Pkts/Bytes)  FP-Up(Pkts/Bytes)
-------------------- ---------- ---------- ---------- ---------- ---------- --------------
  -------------------  -----------------

SGACL Match stats:
ACL Name             Pkts-Down  Bytes-Down Pkts-Up    Bytes-Up   Pkts dropped
-------------------- ---------- ---------- ---------- ---------- -------------
ACL2611                      4       1432          0          0             2
ACL2612                      4       4112          0          0             2
ACL2621                      8       5704          0          0             4
....
```

```
Total subscribers matching specified criteria: 1
```

show subscribers user-plane-only callid *callid_value* flows full

The output of this CLI command is enhanced with the **Uplink SGT**, **Downlink SGT**, and **Matched SGACL** fields for displaying information related to User Access Control through SGACL.

```
show subscribers user-plane-only callid 00004e21 flows full
  Callid: 00004e21
  Interface Type: Sxab
  IP address: n/a

  Flow ID: 1:1
  Uplink pkts: 1                    Downlink pkts: 1
  Uplink bytes: 1040                Downlink bytes: 40
…
  Downlink Sfp Id: NA
  Uplink SGT: 0xA1
  Downlink SGT: 0xB5
  Matched SGACL: ACL_123
```

show user-plane-service sgt-column summary

The output of this CLI command shows the summary of all the available D-SGT values on UPF along with their refresh timers as received from the ISE server.

The following example output shows the D-SGT values of **D-SGT Columns fetched** and **D-SGT Version** fields.

```
show user-plane-service sgt-column summary
 D-SGT Columns fetched:
 D-SGT      Version
 ---------------------------------------
 65535         0
   7           0
 Total D-SGT column(s) found: 2
```

show user-plane-service sgt-column dsgt

The output of this CLI command shows the summary of the SGACL matrix per D-SGT with the respective SGACL mapping per D-SGT and S-SGT, which are received from ISE during the D-SGT query.

The following example output shows the **D-SGT**, **Refresh TimeS-SGT**, **SGACL Name**, **Version**, and **Total SGT column(s) found** fields.

```
show user-plane-service sgt-column dsgt 26

    D-SGT: 26            Refresh Time: 86400
    S-SGT: 1
    SGACL Name           Version
    ---------------------------------------
    ACL2611                              0
    ACL2612                              0

    S-SGT: 2
    SGACL Name                      Version
    ---------------------------------------
    ACL2621                              0

  Total SGT column(s) found: 1
```

show user-plane-service sgacl name

The output of this CLI command shows the specific SGACLs definition rule lines as received from ISE during the SGACL query.

```
show user-plane-service sgacl name AACL2
 SGACL Name: AACL2
 permit ip
 Total SGACL(s) found: 1
```

show user-plane-service statistics sgacl all

The output of this CLI command shows the packet match statistics as per SGACL.

```
show user-plane-service statistics sgacl all

ACL Name              Pkts-Down  Bytes-Down Pkts-Up    Bytes-Up   Pkts dropped
--------------------- ---------- ---------- ---------- ---------- -------------
SGACL1-REFRESH2          6           1156       0          0          2
SGACL1-REFRESH1         13           9344       0          0          6
Allow All                0              0       0          0          0

Total SGACL(s) : 3
```

show user-plane-service statistics drop-counter

The output of this CLI command is enhanced to shows the dropped packets due to the SGACL application.

```
show user-plane-service statistics drop-counter
Packet Drop Data Statistics:
        ------------------------------------------------
...
        FastPath Misc Drops:
            Overload Protection:                    0
            Invalid Client:                         0
            Stream ID 0:                            0
            Invalid Stream ID:                      0
        OHR Mismatch Packet Drops:                  0
       SGACL Packet Drops:                      8
        SGACL No Policy Packet Drops:              2
        No Default SGT cell Packet Drops:          0
```

Note

- **No Default SGT cell Packet Drops** is an obsolete counter. The existing design allows packets to be passed only if no default ACL is available.
- The statistics for ISE server REST API request and response are supported.

show apimgr statistics ise-server

The **show apimgr statistics ise-server** displays the status for ISE servers statistics.

```
[local]qvpc-si# show apimgr statistics ise-server

Ise-Server Connection Statistics:

Server: 10.10.10.30           Current state: Active
Last Inactive Time:          0000:00:00:00:00:00
```

```
Request sent:                           5
Response Success:                       4
Response Fail:                          1

Server: 10.10.10.40           Current state: Inactive
Last Inactive Time:            2025:01:16:12:30:11
Request sent:                           0
Response Success:                       0
Response Fail:                          0
```

SNMP trap

The trap **ISEServerSwitchoverOccured** is generated every time the UPF switches to another ISE server.

OAM Support

Statistics

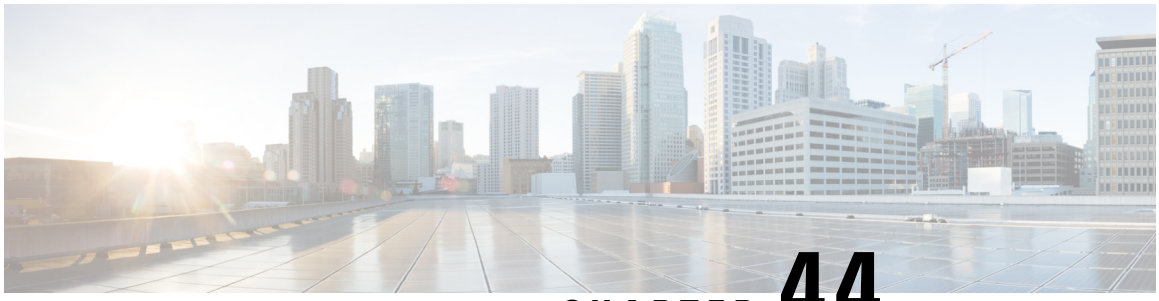
Following new bulk statistics are supported for the user access control through SGACL feature.

SCHEMA: UPF	
Statistics	**Description**
show user-plane-service statistics sgacl all	The count of the packet match statistics as per SGACL.
show user-plane-service statistics drop-counter	The count of the dropped packets due to the SGACL application.
show apimgr statistics ise-server	The count of UPFs API interaction with ISE server.

Bulk Statistics

Following new bulk statistics are supported for the user access control through SGACL feature.

SCHEMA: UPF	
Statistics	**Description**
downlink-total-pkts-sgacl-matched	Total downlink packets matched against SGACL
downlink-total-bytes-sgacl-matched	Total downlink bytes matched against SGACL
downlink-total-pkts-sgacl-dropped	Total downlink packets dropped due to SGACL match
downlink-total-bytes-sgacl-dropped	Total downlink bytes dropped due to SGACL match



CHAPTER 44

Session Recovery

- Feature Summary and Revision History, on page 483
- Feature Description, on page 483
- How it Works, on page 484
- Configuring the System to Support Session Recovery, on page 484

Feature Summary and Revision History

Summary Data

Applicable Product(s) or Functional Area	5G-UPF
Applicable Platform(s)	VPC-SI SMI
Feature Default Setting	Disabled – License Required
Related Changes in this Release	Not Applicable
Related Documentation	Not Applicable

Revision History

Revision Details	Release
First introduced.	2020.02.0

Feature Description

With robust hardware failover and redundancy protection, any hardware or software failures on the system can quickly be corrected. However, software failures can occur for numerous reasons, often without prior indication.

This chapter describes the Session Recovery feature that provides seamless failover and reconstruction of subscriber session information in the event of a hardware or software fault.

Important Session Recovery is a licensed Cisco feature. A separate feature license may be required. Contact your Cisco Account representative for detailed information on specific licensing requirements.

How it Works

This section provides an overview of how this feature is implemented and the recovery process.

The Session Recovery feature provides seamless failover and reconstruction of subscriber session information in the event of a hardware or software fault within the system preventing a fully connected user session from being disconnected.

Session recovery is performed by mirroring key software processes (for example, session manager and AAA manager) within the system. These mirrored processes remain in an idle state (standby-mode) wherein they perform no processing, until they may be needed in the event of a software failure (for example, a session manager task aborts).

There are some situations wherein session recovery may not operate properly. More software or hardware failures occur during the session recovery operation. For example, an AAA manager fails while the state information it contained was being used to populate the newly activated session manager task.

Important After a session recovery operation, some statistics, such as those collected and maintained on a per manager basis (AAA Manager, Session Manager, and so on) are in general not recovered, only accounting and billing related information is checkpointed and recovered.

Configuring the System to Support Session Recovery

The following procedures allow you to configure the session recovery feature for either an operational system that is currently in-service (able to accept incoming calls) or a system that is out-of-service (not part of your production network and, therefore, not processing any live subscriber/customer data).

Important The session recovery feature, even when the feature use key is present, is disabled by default on the system.

Enabling Session Recovery

As noted earlier, session recovery can be enabled on a system that is out-of-service (OOS) and does not yet have any contexts configured, or on an in-service system that is currently capable of processing calls. However, if the system is in-service, it must be restarted before the session recovery feature takes effect.

Enabling Session Recovery on an Out-of-Service System

The following procedure is for a system that does not have any contexts configured.

To enable the session recovery feature on an Out-of-Service system, perform the following procedure. This procedure assumes that you begin at the EXEC mode prompt.

Procedure

Step 1 At the EXEC mode prompt, verify that the session recovery feature is enabled through the session and feature use licenses on the system by running the **show license info** command.

If the current status of the Session Recovery feature is Disabled, you cannot enable this feature until a license key is installed in the system.

Step 2 Use the following configuration example to enable session recovery.

```
configure
   require session recovery
   end
```

Note
After you configure this command, you must save the configuration and then reload the chassis for the command to take effect. For information on saving the configuration file and reloading the chassis, refer to the System Administration Guide for your deployment.

Step 3 Save your configuration as described in *Verifying and Saving Your Configuration*.

The system, when started, enables session recovery, creates all mirrored "standby-mode" tasks, and performs packet processing card reservations and other operations automatically.

Step 4 After the system has been configured and placed in-service, you must verify the preparedness of the system to support this feature as described in *Viewing Session Recovery Status* section.

Enabling Session Recovery on an In-Service System

When enabling session recovery on a system that already has a saved configuration, the session recovery commands are automatically placed before any service configuration commands in the configuration file.

To enable the session recovery feature on an in-service system, perform the following procedure. This procedure assumes that you begin at the EXEC mode prompt.

Procedure

Step 1 At the EXEC mode prompt, verify that the session recovery feature is enabled through the session and feature use licenses on the system by running the **show license info** command:

If the current status of the Session Recovery feature is Disabled, You cannot enable this feature until a license key is installed in the system.

Step 2 Use the following configuration example to enable session recovery.

```
configure
  require session recovery
  end
```

This feature does not take effect until after the system has been restarted.

Step 3 Save your configuration as described in *Verifying and Saving Your Configuration*.

Step 4 Perform a system restart by entering the **reload** command:

The following prompt appears:

```
Are you sure&quest; [Yes|No]:
```

Confirm your desire to perform a system restart by entering **yes**.

The system, when restarted, enables session recovery and creates all mirrored "standby-mode" tasks, performs packet processing card reservations, and other operations automatically.

Step 5 After the system has been restarted, you must verify the preparedness of the system to support this feature as described in *Viewing Session Recovery Status* section.

More advanced users may opt to simply insert the **require session recovery** command syntax into an existing configuration file using a text editor or other means, and then applying the configuration file manually. Exercise caution when doing this to ensure that this command is placed among the first few lines of any existing configuration file; it must appear before the creation of any nonlocal context.

Disabling the Session Recovery Feature

To disable the session recovery feature on a system, enter the **no require session recovery** command from the Global Configuration mode prompt.

Important If this command is issued on an in-service system, then the system must be restarted by issuing the **reload** command.

Viewing Session Recovery Status

To determine if the system is capable of performing session recovery, when enabled, enter the **show session recovery status verbose** command from the Exec mode prompt.

The output of this command should be similar to the examples shown below.

```
[local]host_name# show session recovery status
Session Recovery Status:
     Overall Status             :  SESSMGR Not Ready For Recovery
     Last Status Update         :  1 second ago

[local]host_name# show session recovery status
Session Recovery Status:
    Overall Status              :  Ready For Recovery
    Last Status Update          :  8 seconds ago

[local]host_name# show session recovery status verbose
```

```
Session Recovery Status:
    Overall Status              :  Ready For Recovery
    Last Status Update          :  2 seconds ago

              ----sessmgr----       ----aaamgr----         demux
 cpu state    active    standby     active    standby      active    status
---- -------  ------    -------     ------    -------      ------    ------------
 1/0 Active   7         1           7         1            7         Good
[local]host_name#
```

Viewing Recreated Session Information

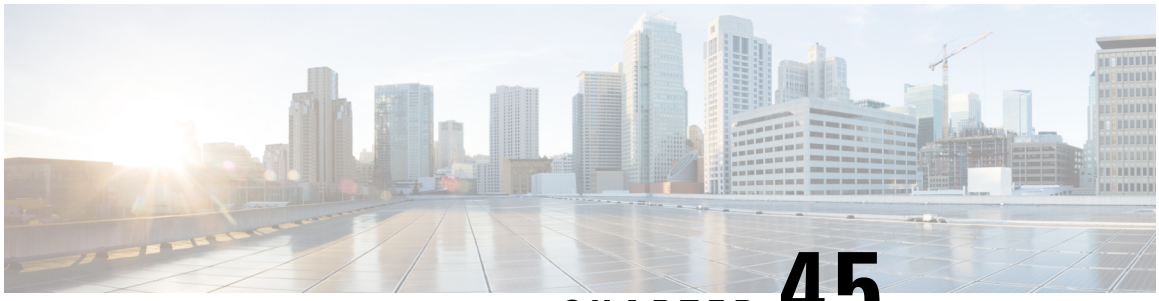
To view session state information and any session recreation status, enter the following command:

show subscriber debug-info callid *id*

The following example shows the output of this command both before and after a session recovery operation has been performed. The "Redundancy Status" fields in this example have been bold-faced for clarity.

```
username: user1              callid: 01ca11b1                  msid: 0000100003
    Card/Cpu: 4/2
    Sessmgr Instance: 7
    Primary callline:
   Redundancy Status: Original Session
      Checkpoints        Attempts        Success        Last-Attempt       Last-Success
          Full:               69             68             29800ms            29800ms
          Micro:             206            206             20100ms            20100ms
      Current state: SMGR_STATE_CONNECTED
      FSM Event trace:
            State                              Event
            SMGR_STATE_OPEN                    SMGR_EVT_NEWCALL
            SMGR_STATE_NEWCALL_ARRIVED         SMGR_EVT_ANSWER_CALL
            SMGR_STATE_NEWCALL_ANSWERED        SMGR_EVT_LINE_CONNECTED
            SMGR_STATE_LINE_CONNECTED          SMGR_EVT_LINK_CONTROL_UP
            SMGR_STATE_LINE_CONNECTED          SMGR_EVT_AUTH_REQ
            SMGR_STATE_LINE_CONNECTED          SMGR_EVT_IPADDR_ALLOC_SUCCESS
            SMGR_STATE_LINE_CONNECTED          SMGR_EVT_AUTH_SUCCESS
            SMGR_STATE_LINE_CONNECTED          SMGR_EVT_UPDATE_SESS_CONFIG
            SMGR_STATE_LINE_CONNECTED          SMGR_EVT_LOWER_LAYER_UP
    Data Reorder statistics
    Total timer expiry:                0        Total flush (tmr expiry):    0
            Total no buffers:          0        Total flush (no buffers):    0
            Total flush (queue full):  0        Total flush (out of range):  0
            Total flush (svc change):  0        Total out-of-seq pkt drop:   0
            Total out-of-seq arrived:  0
    IPv4 Reassembly Statistics:
              Success:                 0        In Progress: 0
              Failure (timeout):       0        Failure (no buffers): 0
              Failure (other reasons): 0
    Redirected Session Entries:              Allowed:
2000            Current:                                    0
                Added:                                      0                Deleted:
                                  0
                Revoked for use by different subscriber: 0
    Peer callline:
     Redundancy Status: Recreated Session
    Checkpoints        Attempts        Success        Last-Attempt        Last-Success
          Full:               0              0                 0ms                 0ms
          Micro:              0              0                 0ms                 0ms
      Current state: SMGR_STATE_CONNECTED
      FSM Event trace:
            State                                                    Event
```

```
        SMGR_STATE_LINE_CONNECTED           SMGR_EVT_LOWER_LAYER_UP
        SMGR_STATE_CONNECTED                SMGR_EVT_AUTH_REQ
        SMGR_STATE_CONNECTED                SMGR_EVT_AUTH_SUCCESS
        SMGR_STATE_CONNECTED                SMGR_EVT_REQ_SUB_SESSION
        SMGR_STATE_CONNECTED                SMGR_EVT_RSP_SUB_SESSION
        SMGR_STATE_CONNECTED                SMGR_EVT_ADD_SUB_SESSION
        SMGR_STATE_CONNECTED                SMGR_EVT_AUTH_REQ
        SMGR_STATE_CONNECTED                SMGR_EVT_AUTH_SUCCESS
        SMGR_STATE_CONNECTED                SMGR_EVT_AUTH_REQ
        SMGR_STATE_CONNECTED                SMGR_EVT_AUTH_SUCCESS
        SMGR_STATE_CONNECTED                SMGR_EVT_AUTH_REQ
        SMGR_STATE_CONNECTED                SMGR_EVT_AUTH_SUCCESS
        SMGR_STATE_CONNECTED                SMGR_EVT_AUTH_REQ
        SMGR_STATE_CONNECTED                SMGR_EVT_AUTH_SUCCESS
        SMGR_STATE_CONNECTED                SMGR_EVT_AUTH_REQ
        SMGR_STATE_CONNECTED                SMGR_EVT_AUTH_SUCCESS
Data Reorder statistics
        Total timer expiry:        0         Total flush (tmr expiry):  0
        Total no buffers:          0         Total flush (no buffers):  0
        Total flush (queue full):  0         Total flush (out of range):0
        Total flush (svc change):  0         Total out-of-seq pkt drop: 0
        Total out-of-seq arrived:  0
IPv4 Reassembly Statistics:
          Success:                   0       In Progress:               0
          Failure (timeout):         0       Failure (no buffers):      0
          Failure (other reasons):   0
Redirected Session Entries:
          Allowed:               2000         Current:             0
          Added:                              Deleted:             0
          Revoked for use by different subscriber:  0
```



CHAPTER 45

Session Report Rejection Procedure

- Feature Summary and Revision History, on page 489
- Feature Description, on page 489
- OAM Support, on page 491

Feature Summary and Revision History

Summary Data

Applicable Product(s) or Functional Area	5G-UPF
Applicable Platform(s)	VPC-SI SMI
Feature Default Setting	Enabled – Always-on
Related Changes in this Release	Not Applicable
Related Documentation	*UCC 5G SMF Configuration and Administration Guide*

Revision History

Revision Details	Release
First introduced.	2021.02.0

Feature Description

The SMF rejects the UPF-originated Session Report Request with a specific cause code during any mismatch in the charging configuration of SMF and UPF.

For any session report rejection by the SMF, the UPF locally purges the sessions. The SMF is unaware of the purging operation and continues to send the N4 message to the UPF. This action triggers the UPF to send “context not found” message to the SMF for the locally purged sessions.

This behavior impacts the UE experience and results in the loss of charging data. So, the current implementation of handling the session report errors is modified to avoid local purging of sessions on the UPF and also to support graceful clearing of sessions.

With this modification, the UPF ignores the Session Report Error Response. The SMF triggers the Session Deletion Request followed by the rejection of Session Report. The UPF responds to the delete request and clears the session gracefully.

Relationships to Other Features

This feature involves implementing some behavioral changes to the SMF and the UPF. The new CLI configuration in SMF aids in controlling this behavior. For details on the SMF behavioral changes, see the *Support for Session Report Rejection Procedure* chapter in the *UCC 5G SMF Configuration and Administration Guide*.

Call Flow

The following figure explains the flow of Session Report rejection prcedure.

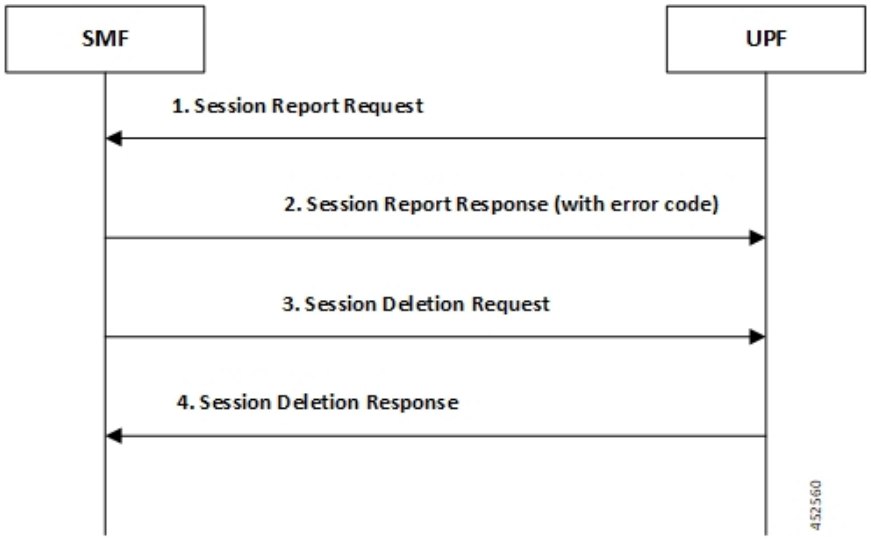


Table 106: Session Report Rejection Procedure Call Flow

Step	Description
1	UPF sends Session Report message to SMF. (SX_SESSION_REPORT_REQUEST)
2	SMF rejects the message with some error code. (SX_SESSION_REPORT_RESPONSE with error code)
3	UPF ignores the Session Report error response. SMF initiates Session Delete Request to tear down the session gracefully. (SX_SESSION_DELETION_REQUEST)

Step	Description
4	UPF responds to the delete request and clears the session gracefully. (SX_SESSION_DELETION_RESPONSE)

OAM Support

This section describes operations, administration, and maintenance information for this feature.

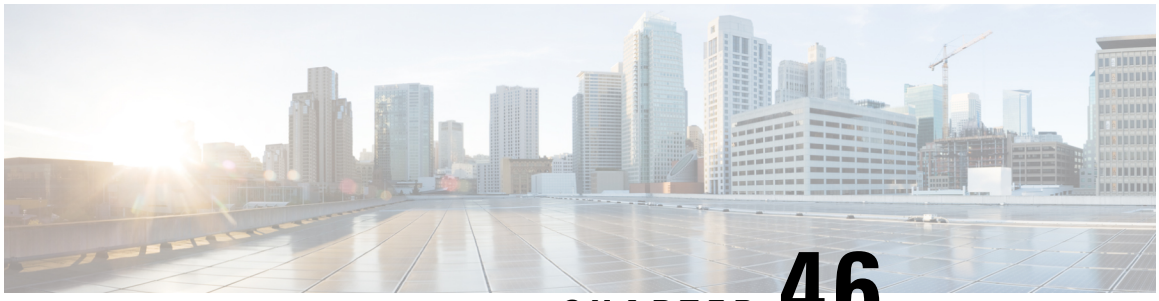
Show Command(s) and/or Output(s)

show user-plane-service statistics all

The counter, "Skipped Local Purge", is added to the **show user-plane-service statistics all** CLI command under the section “PDNs Released By Reason”.

- **Skipped Local Purge**: Increments whenever Session Report error is triggered from SMF/CP to UPF/UP.

Show Command(s) and/or Output(s)



CHAPTER 46

Smart Licensing

- Feature Summary and Revision History, on page 493
- Overview, on page 494
- Configuring Smart Licensing, on page 499
- Monitoring and Troubleshooting Smart Licensing, on page 500

Feature Summary and Revision History

Summary Data

Table 107: Summary Data

Applicable Product (s) or Functional Area	5G-UPF
Applicable Platforms	VPC-SI SMI
Feature Default Setting	Disabled – Configuration Required
Related Changes in this Release	Not Applicable
Related Documentation	*UCC 5G UPF Configuration and Administration Guide*

Revision History

Revision Details	Release
First introduced.	2021.01.0

Overview

Ultra Cloud Core 5G User Plane Function (UPF) supports Smart Licensing. Smart Licensing is a cloud-based approach to licensing that simplifies the purchase, deployment, and management of Cisco software assets. Entitlements are purchased through your Cisco account via Cisco Commerce Workspace (CCW) and immediately deposited into your Virtual Account for usage. This eliminates the need to install license files on every device. Products that are smart-enabled, communicate directly to Cisco to report consumption. A single location is available to customers to manage Cisco software licenses - the Cisco Smart Software Manager (CSSM). License ownership and consumption are readily available to help make better purchase decision based on consumption or business need.

See https://www.cisco.com/c/en/us/buy/smart-accounts/software-licensing.html for more information about Cisco Smart Licensing.

Comparison Between Legacy Licensing and Smart Licensing

Cisco employs two types of license models - Legacy Licensing and Smart Software Licensing. **Legacy Licensing** consists of software activation by installing Product Activation Keys (PAK) on to the Cisco product. A Product Activation Key is a purchasable item, ordered in the same manner as other Cisco equipment and used to obtain license files for feature set on Cisco Products. **Smart Software Licensing** is a cloud-based licensing of the end-to-end platform leveraging few tools that authorize and deliver license reporting. Smart Software Licensing functionality incorporated into Network Functions (NFs) completes the product registration, authorization resulting in reporting services available to the end customer.

Evaluation Period

A 90-day evaluation period is granted for all licenses in use. During this period, feature licenses can be used without limitation, and up to one counting license each can be used. The evaluation period ends when the system registers successfully with the CSSM or Cisco.com. Licensed functionality is blocked when this 90-day period expires.

UPF performs license enforcement for on/off feature licenses. Each on/off feature license is tied to service licenses, which potentially use those on/off features. When an Out of Compliance (OOC) is detected for an on/off license, new calls for the corresponding services are dropped, subject to the following conditions:

- Each on/off feature license is given a 90-day grace (evaluation) period. During this period, the system generates SNMP traps to inform of the unavailability of valid licenses. To resolve the OOC, corrective action is needed such as purchasing and registering licenses for this feature, or disabling the feature.
- If the feature is still OOC after the 90-day grace period, UPF enforces the OOC state based on a predefined policy for each license. If enforcement is required, new calls for the services corresponding to the on/off licenses are dropped.

The following CLI commands can be used to display details about the enforcement of Smart Licenses in use:

```
show license enforcement policy
show license enforcement status [ allowed | blocked ] [ feature | service
 ]
```

Cisco Smart Software Manager

Cisco Smart Software Manager (CSSM) enables the management of software licenses and Smart Account from a single portal. The interface allows you to activate your product, manage entitlements, and renew and upgrade software. A functioning Smart Account is required to complete the registration process. To access the Cisco Smart Software Manager, see https://software.cisco.com.

Smart Accounts/Virtual Accounts

A Smart Account provides a single location for all Smart-enabled products and entitlements. It helps speed procurement, deployment, and maintenance of Cisco Software. When creating a Smart Account, you must have the authority to represent the requesting organization. After submitting, the request goes through a brief approval process.

A Virtual Account exists as a sub-account withing the Smart Account. Virtual Accounts are a customer-defined structure based on organizational layout, business function, geography or any defined hierarchy. They are created and maintained by the Smart Account administrator.

See https://software.cisco.com to learn about, set up, or manage Smart Accounts.

Smart Licensing Mode

The Smart Licensing Mode is categorized as follows:

- **Reporting Licenses (Parent Licenses)**: The Parent Licenses are reported to backend license server (CSSM) and accounted for usage of licenses. For each Parent Licenses, the entitlement tags are created and the same is used to identify the type service or feature.
- **Non-Reporting Licenses (Child Licenses)**: The Child Licenses are not reported to backend license server (CSSM) and these licenses are enabled by default with the Parent Licenses. For Child Licenses, the entitlement tags are not created.

That is to say, Smart License enables all Parent and Child Licenses based on the Product Type that is configured. However, the reporting is done only for Parent Licenses.

The state of Smart Licensing Agent is persistent across reboot and crashes.

Request a Cisco Smart Account

A Cisco Smart Account is an account where all products enabled for Smart Licensing are deposited. A Cisco Smart Account allows you to manage and activate your licenses to devices, monitor license use, and track Cisco license purchases. Through transparent access, you have a real-time view into your Smart Licensing products. IT administrators can manage licenses and account users within your organization's Smart Account through the Smart Software Manager.

Procedure

Step 1 In a browser window, enter the following URL:

```
https://software.cisco.com
```

Step 2 Log in using your credentials, and then click **Request a Smart Account** in the **Administration** area.

The **Smart Account Request** window is displayed.

Step 3 Under **Create Account**, select one of the following options:

- **Yes, I have authority to represent my company and want to create the Smart Account** – If you select this option, you agree to authorization to create and manage product and service entitlements, users, and roles on behalf of your organization.
- **No, the person specified below will create the account** – If you select this option, you must enter the email address of the person who will create the Smart Account.

Step 4 Under **Account Information**:

a) Click **Edit** beside **Account Domain Identifier**.

b) In the **Edit Account Identifier** dialog box, enter the domain, and click **OK**. By default, the domain is based on the email address of the person creating the account and must belong to the company that will own this account.

c) Enter the **Account Name** (typically, the company name).

Step 5 Click **Continue**.

The Smart Account request will be in pending status until it has been approved by the Account Domain Identifier. After approval, you will receive an email confirmation with instructions for completing the setup process.

Software Tags and Entitlement Tags

Tags for the following software and entitlements have been created to identify, report, and enforce licenses.

Software Tags

Software tags uniquely identify each licenseable software product or product suite on a device. The following software tags exist for UPF.

Product Type / Description	Software Tag
UPF Ultra Cloud Core - User Plane Function (UPF)	regid.2020-04.com.cisco.UPF, 1.0_bc18a9ff-e0ea-4476-a250-04ebf7839c4c

Reporting (Parent) Entitlement Tags for UPF

The following entitlement tags indentify licenses in use for UPF.

License Display Name/Description	Entitlement Tag	Tag Name
UCC 5G UPF Base Lic Ultra Cloud Core - User Plane Function (UPF), Base Minimum	regid.2020-08.com.cisco.F_UPF_BASE, 1.0_776395f3-8b8d-46e1-ac6e-0bd2306ef3b6	F_UPF_BASE

License Display Name/Description	Entitlement Tag	Tag Name
UCC 5G UPF Instance Ultra Cloud Core - User Plane Function (UPF) Instance	regid.2020-08.com.cisco.F_UPF_INS, 1.0_5cd68c07-152a-48c6-b143-4dc60eb111e5	F_UPF_INS
UCC 5G UPF Content Filtering UCC - User Plane Function (UPF) Content Filtering	regid.2020-08.com.cisco.F_UPF_CF, 1.0_db1ad035-0753-40e3-9f60-2a26135c3569	F_UPF_CF
UCC 5G UPF CUTO 1K UCC - User Plane Function (UPF) Cisco Ultra Traffic Optimization (CUTO) 1K Sessions	regid.2020-08.com.cisco.L_UPF_CUTO_1K, 1.0_41ee8e5e-4307-4299-a77c-c15bbb792e2c	L_UPF_CUTO_1K
UCC 5G UPF 1K Sess Ultra Cloud Core - User Plane Function (UPF), 1K Sessions	regid.2020-08.com.cisco.L_UPF_SAE_1K, 1.0_5d16e2f6-808a-45ff-8691-f215d5ba2bea	L_UPF_SAE_1K

Non-reporting (Child) License List

In this release, the following Child Licenses are enabled by default when the Parent Licenses are enabled.

License Description	License Type
PGW 1k Sessions	Counting
SGW 1k Sessions	Counting
GGSN 1k Sessions	Counting
Per Subscriber Stateful Firewall 1k Sessions	Counting
ENAT 1k Sessions	Counting
Enhanced Charging Bundle 1	Counting
Enhanced charging bundle 2	On/Off
Dynamic policy interface	On/Off
Enhanced LI service	On/Off
Lawful intercept	On/Off
Session recover	On/Off
Radius AAA server group	On/Off
IPv6	On/Off

License Description	License Type
Intelligent Traffic Control	On/Off
DIAMETER Closed-Loop Charging Interface	On/Off
Per-Subscriber Traffic Policing/Shaping	On/Off
Dynamic Radius extensions (CoA and PoD)	On/Off
Proxy MIP	On/Off
FA	On/Off
IPSec	On/Off
Inter-Chassis Session Recovery	On/Off
ICSR/SR Performance Improvements	On/Off
ICSR Enhanced Recovery for Data and Control Plane, 1K Sessions	On/Off
MPLS	On/Off
TACACS+	On/Off
NAT/PAT With DPI	On/Off
Rate Limiting Function (Throttling)	On/Off
Overcharging Protection for EPC-GW	On/Off
Overcharging Protection Upgrade for EPC-GW	On/Off
ADC Trigger Over Gx, 1K Sessions	On/Off
Gx Based Virtual APN Selection, 1K Sessions	On/Off
EPC-GW Support for Wi-Fi Integration, 1K Sessions	On/Off
EPC-GW Non-Standard QCI Support, 1K Sessions	On/Off
Local Policy Decision Engine	On/Off
Header Enrichment	On/Off
HTTP Header Encryption	On/Off
HTTP Header Enrichment and Encryption	On/Off
Broadcast & Multicast Services	On/Off
Integrated Content Filtering Provisioned Service	On/Off
Application Detection and Control 1k Sessions	Counting
5G NSA Feature Set 100K Sess VPCSW Active 1k Sessions	Counting
5G NSA Enablement Fee, Network Wide	On/Off
Multimedia Priority Service Feature Set,1K Sessions	On/Off
EPC Gw VoLTE enhancements	On/Off
DNS Snooping	On/Off

Configuring Smart Licensing

Before you begin, ensure you have:

- Created a Smart Licensing account on https://software.cisco.com.
- Registered your products on https://software.cisco.com using the Product Instance Registration tokens created as part of a Smart Account or Virtual Account.
- Enabled a communication path between the UPF system to the CSSM server or Cisco.com.

Enable Smart Licensing

By default, Smart Licensing is disabled in UPF. To enable Smart Licensing, enter the following Global Configuration mode commands:

```
configure
  license smart product upf
  license smart enable
  end
```

NOTE: Before enabling Smart Licensing, Product Type must be configured to enable default licenses that are based on product type.

Enter the following command to verify the configuration:

```
show configuration | grep license
```

Register the Device with Cisco

Using the Product Instance Registration token ID provided when you registered the products on https://software.cisco.com, register the system using the following EXEC mode command:

```
license smart register idtoken token
```

The system now automatically reports an entitlement usage count to the CSSM server and receives a compliance status. This also removes the system from "Evaluation Mode".

To show the compliance status, enter any of the following EXEC mode commands:

```
show license status
show license summary
show license statistics
```

The registration for the system is renewed automatically every 180 days. If needed, use the following EXEC mode command to renew the registration information manually:

```
license smart renew id
```

The license authorization for the system is renewed automatically every 30 days. If needed, use the following EXEC mode command to renew the license authorization manually:

```
license smart renew auth
```

To unregister a device, enter the following EXEC mode command:

```
license smart deregister
```

Changing Smart Transport URL

Smart Agent uses Smart Transport to communicate to Cisco CSSM server. Smart Transport uses the configured URL to identify the destination URL where CSSM is reachable. This will not initiate any communication with Cisco. If needed, enter the following configuration mode commands:

```
configure
   license smart transport smart
   license smart url https_link
```

Handling Out of Compliance

If there are not enough licenses in the virtual account for a given SKU, CSSM sends an Out Of Compliance (OOC) message to the device. The system stops allowing extra sessions until the OOC state is cleared. The OOC state is cleared when the device receives an authorized response.

Monitoring and Troubleshooting Smart Licensing

Enter the following EXEC mode command to verify the Smart Licensing configuration:

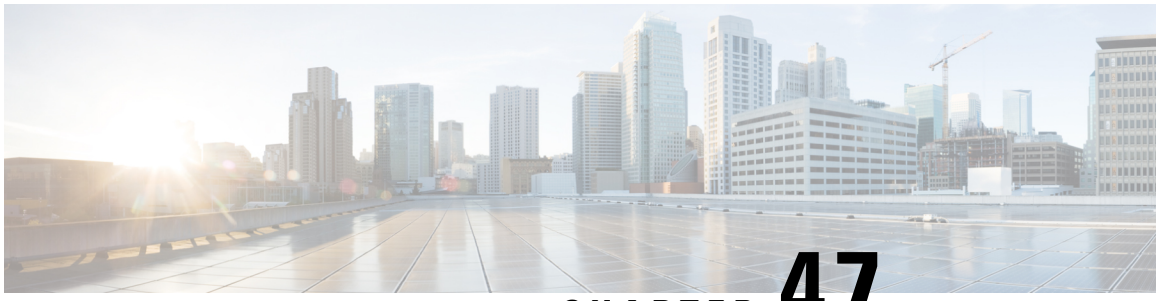
```
show configuration | grep license
```

The following EXEC mode commands display information about Smart Licensing:

```
show license { all | enforcement | smart-tags | statistics | status |
summary | tech-support | udi | usage }
```

NOTES:

- **all** - Shows a superset of information that includes show status, show usage, show UDI, as well as the Smart Licensing agent version.
- **enforcement { policy | status [allowed | blocked] [feature | service] }** - Shows the enforcement policy applied or current enforcement status of Smart Licenses. Status information can be filtered to show only the licenses which are currently allowed or blocked, or by type (feature license or service license).
- **smart-tags [feature | service]** - Shows the features and services that are currently supported and the corresponding Smart Entitlement Tag.
- **statistics [verbose]** - Shows individual feature license status.
- **status** - Shows overall Smart Licensing status information.
- **summary** - Shows summary of Smart Licensing status.
- **tech-support** - Shows information useful for debugging issues with Smart Licensing.
- **udi** - Shows details for all Unique Device Identifiers (UDI).
- **usage** - Shows the usage information for all entitlements that are currently in use.



CHAPTER 47

Software Management Operations

- Feature Summary and Revision History, on page 501
- Overview, on page 502
- SNMP Traps, on page 503
- Limitations, on page 503
- Health Checks, on page 503
- Build Upgrade, on page 505
- UPF Upgrade, on page 507
- UPF Downgrade, on page 507

Feature Summary and Revision History

Summary Data

Table 108: Summary Data

Applicable Product(s) or Functional Area	5G-UPF
Applicable Platform(s)	VPC-SI SMI
Feature Default Setting	Enabled – Always-on
Related Changes in This Release:	Not Applicable
Related Documentation	Not Applicable

Revision History

Table 109: Revision History

Revision Details	Release
First introduced.	2022.01.2

Overview

5G UPF supports backward compatibility of software releases on the SMF and the UPF. The feature allows seamless upgrade/downgrade of the software from/to one previous release (N-1). The functionality includes support for the following:

- N-1 compatibility of software releases on two UPFs in ICSR mode—allows seamless upgrade of UPFs from one version to another in UPF 1:1 redundancy scenario.
- N-1 compatibility of software releases between SMF and UPF—allows seamless upgrade of the associated SMF or UPF from one version to another.
- N-1 compatibility of software releases between SMF and UPF with multi-Sx—allows seamless upgrade of the associated SMF or UPF from one version to another in multi-Sx scenario.



Important Contact your Cisco Account representative for procedural assistance prior to upgrading or downgrading your software versions.

Version Exchange between SMF and UPF

Version/release information is exchanged when SMF and UPF pairs. The release information exchange also occurs when the UPF pairs with a Standby UPF (in 1:1 redundancy scenario) through the heart beat message exchanged between Active and Standby.

When incompatible releases are paired, an Alarm (SNMP trap) is raised. For details, see SNMP Traps section.

To indicate the peer version during the exchange of release information, the following new IE is included in the association request and heartbeat request messages.

Information Elements	P	Condition / Comment								IE Length	IE ID
Peer Version	O	Used to specify the peer GR/PFCP version and StarOS version								4 bytes	245
		Bits									
	Octets	8	7	6	5	4	3	2	1		
	1 to 2	Peer Version IE Type = 245 (decimal)									
	3 to 4	Length = n bytes									
	5 to 8	Peer GR/PFCP Version									
	9 to 12	StarOS GR Version									
	13 to 13	StarOS Version String Length									
	Variable Length	StarOS Version String Value									

SNMP Traps

The following SNMP traps are raised when pairing is done with an incompatible release.

SNMP Trap	Description
SRPPeerUnsupportedVersion	The Active/Standby UPF in higher version raises the SNMP trap when the peer is in a version lower than N-1.
SRPPeerUnsupportedVersionClear	The Active/Standby UPF in higher version raises the SNMP trap to clear the SRPPeerUnsupportedVersion.
SxPeerUnsupportedVersion	The UPF in higher version raises the SNMP trap when the peer is in a version lower than N-1.
SxPeerUnsupportedVersionClear	The UPF in higher version raises an SNMP trap to clear the SxPeerUnsupportedVersion.

Limitations

The following are the known limitations of the feature:

- When the peer version is determined to be lower than the supported N-1 version, the association/pairing is allowed. However, functional aspect of the same isn't guaranteed.

Caution Don't attempt to upgrade from incompatible versions. Contact your Cisco Account representative for the upgrade path and steps.

- Few CLI commands may not be supported in N+1 version.
- The SMF version must be compatible with the UPF version.
- The hardware configuration must be similar in both Active and Standby UPFs.
- SNMP traps are raised by the node on the latest version with respect to the StarOS version. For details, see the SNMP Traps section of this chapter.
- From release 2022.01.2, RCM is checkpoint agnostic to enable support for future UPF releases. Currently RCM supports only N-1 compatibility.

Health Checks

Perform the following health checks after every operation of upgrade, downgrade, or reload of chassis.

1. Check the Service Redundancy Protocol (SRP) information on the Active chassis to avoid issues during an SRP switchover and decide if proactive analysis must be done before the SRP switchover. Use the following CLI commands:
   - **srp validate-configuration**

- **srp validate-switchover**
- **show srp info**

The following is a sample output.

```
Peer Configuration Validation: Complete
Last Peer Configuration Error: None
Last Peer Configuration Event: Wed Mar 18 15:34:02 2022  (1602 seconds ago)
Last Validate Switchover Status: None
Connection State: Connected
```

Check the following parameters:

- **Peer Configuration Validation: Complete**—If it shows "In Progress," you must wait and then execute the **show srp info** CLI command again after 15 seconds (approximately).
- **Last Peer Configuration Error: None**—If you see "Peer Checksum Validation Failure," then it indicates that there are configuration differences between Active and Standby chassis that must be fixed.



Note If you see any error in **Last Peer Configuration Error**, validate the configuration using the **show configuration srp** CLI command on both the Active and Standby UPFs.

- **Last Validate Switchover Status: None**—The output must show as "None." Also, the output must be *Remote Chassis - Ready for Switchover (XX seconds ago)* when the **srp validate-configuration** and **srp validate-switchover** CLI commands are triggered.
- **Connection State: Connected**—The output must show as "Connected."

2. Check subscriber count on both Active and Standby chassis.

   After sessions are up, execute the **show subscribers summary | grep Total** CLI command in the Active chassis. The following is a sample output.

   ```
   show subscribers summary | grep Total
      Total Subscribers:  100
   ```

   On the Standby chassis, execute the **show srp checkpoint statistics | grep allocated** CLI command. The following is a sample output.

   ```
   show srp checkpoint statistics | grep allocated
      Current pre-allocated calls:  100
   ```

3. Check the status of the license by executing the **show license information** CLI command. It must be in "Good (Redundant)" and not in "Expired" state.
4. Check the Session Recovery Status by executing the **show session recovery status verbose** CLI command. The following is a sample output.

   ```
   Session Recovery Status:
     Overall Status        : Ready For Recovery
     Last Status Update    : 7 seconds ago

                 ----sessmgr---  ----aaamgr----  demux
    cpu state    active standby  active standby  active  status
   1/0  Active     8      1        8      1        17     Good
   ```

5. Verify if all the SessMgrs are in Standby-Connected state on the Standby chassis by executing the **show srp checkpoint statistics | grep Sessmgrs** CLI command. The following is a sample output.

```
Number of Sessmgrs:        1
Sessmgrs in Active-Connected state:  0
Sessmgrs in Standby-Connected state: 8
Sessmgrs in Pending-Active state:    0
```

6. Verify the status of all the cards to see if they are in Active or Standby state. The following is a sample output.

```
show card table

Slot          Card Type              Oper State  SPOF  Attach
----------- ---------------------- ----------- ----- ------
1: VC        5-Port Virtual Card     Active        -
```

7. Execute the **show task resources | grep -v good** CLI command, and its output must only display the total number of SessMgrs and sessions.
8. Execute the **show crash list** CLI command to check if there are any new crashes.
9. Execute the **show service all** CLI command to verify that the state is displayed as "Started" and not "Initialized."

Build Upgrade

Backup Configuration

1. Back up the current configuration—Save the current configuration that is used in case of downgrade/upgrade, which probably has all the features and configuration present until now.
2. Collect the **show support details** on both Active and Standby chassis before making any changes or upgrade.
3. Perform Health Checks.

Upgrade Procedure

1. Perform chassis Health Checks on both the nodes.
2. On the secondary chassis (ICSR), which is in Standby state, change boot priority with N+1 build.
3. Reload to the latest build version.
4. Do the new configuration change on Standby chassis (For example, any new CLI, license, or configuration changes).
5. Perform Health Check on the reloaded chassis. Check for any crashes or errors.

Perform Switchover

1. Before SRP switchover from Active to Standby on both chassis, check:
   a. On Active chassis: **show subscriber summary | grep Total**

b. On Standby chassis: **show srp checkpoint statistics | grep allocated**

Note The count must be same for both chassis.

c. On Active and Standby chassis: **show sx peer**

For example:

```
|||||  Sx Service                                                    No of
|||||  ID                                                                 Restart
|||||  |                                                 Recovery              |
Current    Max        Peer
vvvvv  v     Group Name    Node ID         Peer ID     Timestamp             v
Sessions   Sessions   State
----- ---- ----------------- ------------------------- ------------------- ------
--------- --------- --------
CAAXD  22    CPGROUP21   209.165.200.225   50331649    2021-03-17:02:33:55    0
    0          0      NONE

Total Peers:    1
```

Note Peer state must be Active and associated. Peer ID must match on both the chassis.

d. On Standby chassis: **show srp checkpoint statistics | grep Sessmgrs**

Note "Number of Sessmgrs" must be equal to the "Sessmgrs in Standby-Connected state".

e. On Active chassis:

1. **srp validate-configuration**: This CLI command initiates a configuration validation check in the Active chassis. If the validation doesn't have any error, the output of this CLI command is blank.
2. **srp validate-switchover**: Validates if both Active and Standby chassis are ready for a planned SRP switchover. If the chassis is ready for switchover, then the output of this CLI command is blank.
3. **show srp info | grep "Last Validate Switchover Status"** : Output of this CLI command must be as follows.

   ```
   Last Validate Switchover Status: Remote Chassis - Ready for Switchover
   ```

4. **show srp info debug**: Active and Standby chassis must have the same output.

2. On Active chassis: **srp initiate-switchover**
   a. Perform chassis Health Checks on both the nodes. Also check Step 1a and Step 1c under the *Perform Switchover* section. There can be a difference of 5%.
   b. Perform call testing since new sessions are serviced on the new Active chassis.
   c. Upgrade the old Active as mentioned in Step 2 through Step 5 under the *Upgrade Procedure* section.

UPF Upgrade

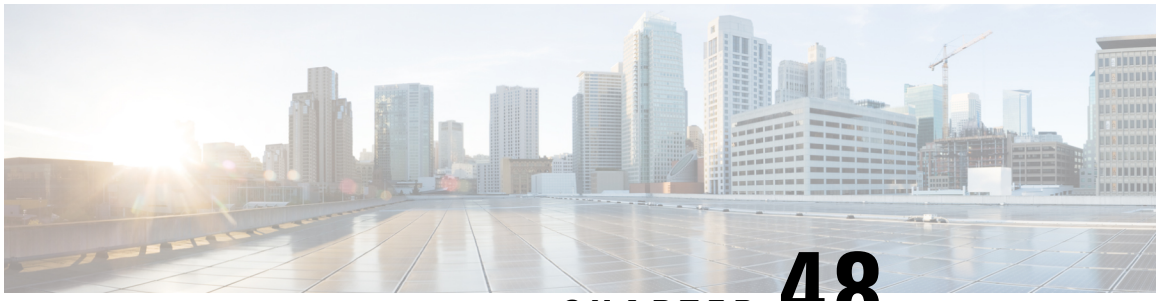
This section describes the procedure for UPF upgrade.

1. Perform Health Check procedure on both the UPF nodes as mentioned in the Health Checks, on page 503 section.
2. Perform Upgrade on Standby UPF as mentioned in the Build Upgrade, on page 505 section.
3. Do "sx-peer configuration" on the upgraded Standby chassis.
4. Perform Health Check on both the UPF nodes, and then do UPF switchover.
5. Upgrade the new Standby UPF as mentioned in the Build Upgrade, on page 505 section.

UPF Downgrade

Perform the following steps to downgrade the UPF:

1. Perform Health Check on the UPF.
2. Change boot priority to the N-1 build on the Standby UPF. Reload the Standby UPF.
3. Do "sx-peer configuration" on the downgraded Standby UPF.
4. Perform Health Check on both the UPF nodes and then do UPF switchover.
5. Perform Step 1 to Step 3 on the new Standby UPF.



CHAPTER 48

Standard QCI Support

- Feature Summary and Revision History, on page 509
- Feature Description, on page 509
- Configurations, on page 510

Feature Summary and Revision History

Summary Data

Table 110: Summary Data

Applicable Product(s) or Functional Area	5G-UPF
Applicable Platform(s)	VPC-SI SMI
Feature Default Setting	Enabled – Always-on
Related Changes in this Release	Not Applicable
Related Documentation	*UCC 5G UPF Configuration and Administration Guide*

Revision History

Revision Details	Release
First introduced.	2021.02.0

Feature Description

The 5G-UPF supports new standard QoS Class Index (QCIs) based on 3GPP TS 23.203 Release 12, for Mission Critical and Push-to-Talk (MC/PTT) applications.

As part of this feature, the following functionalities are supported:

- Establishing a LTE/WiFi/5G-NR call with default bearer QCI/QFI with new standard non-GBR QCIs (69, 70 and 80). These are MC-PTT and 5G-NSA QCIs.
- Establishment of dynamic rule with new standard GBR/non-GBR QCI (65, 66, 69, 70, 80, 82, 83).
- Support for Extended QoS Bit Rates for DCNR-enabled UEs with new standard QCIs (80, 82, and 83).
- LTE to 5G and 5G to LTE HO, WiFi to 5G and 5G to WiFi HO, and LTE to WiFi and WiFi to LTE HO are supported.
- DSCP Marking in UL and DL direction based on new standard QCIs.

Limitations

In this release, predefined rules with the new standard QCIs aren't supported.

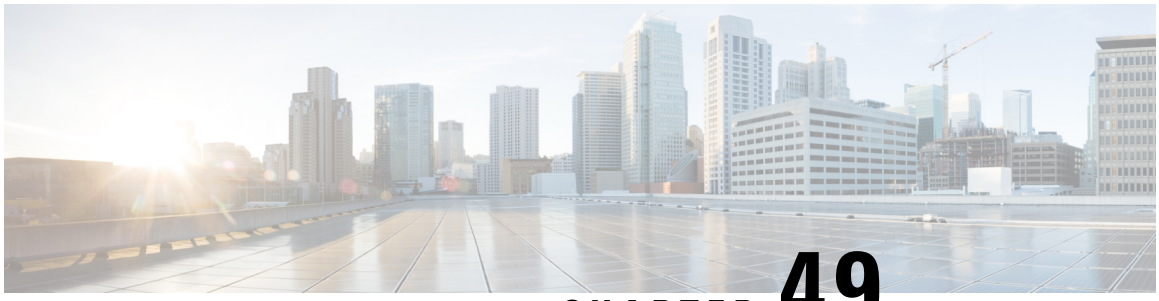
Configurations

There is no configuration (or License) required at SMF or UPF to enable new standard QCIs.

At SMF, the following configuration is required to enable Extended QoS Bit Rates.

```
configure
  profile dnn name
    dcnr { true | false }
    end
```

- **dcnr**: Specifies to enable dual connectivity with new radio (DCNR).



CHAPTER 49

System Logs

This chapter describes how to configure parameters related to the various types of logging and how to viewing their content. It includes the following sections:

- Feature Summary and Revision History, on page 511
- System Log Types, on page 512
- Configuring Event Logging Parameters, on page 513
- Configuring Active Logs, on page 518
- Specifying Facilities, on page 519
- Configuring Trace Logging, on page 528
- Configuring Monitor Logs, on page 528
- Viewing Logging Configuration and Statistics, on page 529
- Viewing Event Logs Using the CLI, on page 529
- Configuring and Viewing Crash Logs, on page 530
- Reducing Excessive Event Logging, on page 532
- Checkpointing Logs, on page 534
- Saving Log Files, on page 534
- Event ID Overview, on page 536

Feature Summary and Revision History

Summary Data

Table 111: Summary Data

Applicable Product(s) or Functional Area	5G-UPF
Applicable Platform(s)	VPC-SI SMI
Feature Default Setting	Disabled – Configuration Required
Related Changes in This Release:	Not Applicable

Related Documentation	*UCC 5G UPF Configuration and Administration Guide*

Revision History

Table 112: Revision History

Revision Details	Release
First introduced.	2020.02.0

System Log Types

There are five types of logs that can be configured and viewed on the system:

Important Not all Event Logs can be configured on all products. Configurability depends on the hardware platform and licenses in use.

- **Event**: Event logging can be used to determine system status and capture important information pertaining to protocols and tasks in use by the system. This is a global function that will be applied to all contexts, sessions, and processes.
- **Active**: Active logs are operator configurable on a CLI instance-by-CLI instance basis. Active logs that are configured by an administrative user in one CLI instance cannot be viewed by an administrative user in a different CLI instance. Each active log can be configured with filter and display properties that are independent of those configured globally for the system. Active logs are displayed in real time as events are generated.
- **Trace**: Trace logging can be used to quickly isolate issues that may arise for a particular connected subscriber session. Traces can be taken for a specific call identification (callid) number, IP address, mobile station identification (MSID) number, or username.
- **Monitor**: Monitor logging records all activity associated with a particular session. This functionality is available in order to comply with law enforcement agency requirements for monitoring capabilities of particular subscribers. Monitors can be performed based on a subscriber's MSID or username.
- **Crash**: Crash logging stores useful information pertaining to system software crashes. This information is useful in determining the cause of the crash.

Important Stateful Firewall and NAT supports logging of various messages on screen if logging is enabled for firewall. These logs provide detailed messages at various levels, like critical, error, warning, and debug. Stateful Firewall and NAT attack logs also provide information on the source IP address, destination IP address, protocol, or attack type for any packet dropped due to an attack and are also sent to a syslog server if configured in the system. For more information on logging support for Stateful Firewall and NAT, see the *Logging Support* chapter of *PSF Administration Guide* or *NAT Administration Guide*.

Configuring Event Logging Parameters

The system can be configured to generate logs based on user-defined filters. The filters specify the facilities (system tasks or protocols) that the system is to monitor and severity levels at which to trigger the generation of the event entries.

Event logs are stored in system memory and can be viewed via the CLI. There are two memory buffers that store event logging information. The first buffer stores the active log information. The second buffer stores inactive logging information. The inactive buffer is used as a temporary repository to allow you to view logs without having data be overwritten. Logs are copied to the inactive buffer only through manual intervention.

Each buffer can store up to 50,000 events. Once these buffers reach their capacity, the oldest information is removed to make room for the newest.

To prevent the loss of log data, the system can be configured to transmit logs to a syslog server over a network interface.

Configuring Event Log Filters

You can filter the contents of event logs at the Exec mode and Global Configuration mode levels.

Exec Mode Filtering

These commands allow you to limit the amount of data contained in logs without changing global logging parameters.

Follow the examples below to filter logs via Exec mode commands.

Active Filtering

logging active [copy runtime filters] [event-verbosity *event_level* **] [pdu-data** *format* **] [pdu-verbosity** *pdu_level* **]**

Notes:

- **copy runtime filters** – Copies the runtime filters and uses that copy to filter the current logging session.
- **event-verbosity** *event_level* – Specifies the level of verboseness to use in logging of events as one of:
  - *min* – Displays minimal information about the event. Information includes event name, facility, event ID, severity level, date, and time.
  - *concise* – Displays detailed information about the event, but does not provide the event source within the system.
  - *full* – Displays detailed information about event, including source information, identifying where within the system the event was generated.
- **pdu-data** *format* – Specifies output format for packet data units when logged as one of:
  - *none* – raw format (unformatted).
  - *hex* – hexadecimal format
  - *hex-ascii* – hexadecimal and ASCII similar to a main-frame dump

- **pdu-verbosity** *pdu_level* – Specifies the level of verboseness to use in logging of packet data units as an integer from 1 through 5, where 5 is the most detailed.

Disable or Enable Filtering by Instance(s)

logging filter active facility *facility* **level** *severity_level* **[critical-info | no-critical-info]**

logging filter { disable | enable } facility *facility* **{ all | instance** *instance_number* **}**

Notes:

- **active** – Indicates that only active processes are to have logging options set.
- **disable** – Disables logging for a specific instance or all instances. This keyword is only supported for aaamgr, hamgr and sessmgr facilities.
- **enable** – Enables logging for a specific instance or all instances. This keyword is only supported for aaamgr, hamgr and sessmgr facilities. By default logging is enabled for all instances of aaamgr, hamgr and sessmgr.
- **facility** *facility* and **level** *severity_level* – Configure the logging filter that determines which system facilities should be logged and at what levels.
- **all | instance** *instance_number* – Specifies whether logging will be disabled or enabled for all instances or a specific instance of aaamgr, hamgr or sessmgr. Run the **show session subsystem facility** *facility* command to identify specific instance numbers.

Note These keywords are only supported with the **disable** and **enable** keywords.

- **level** *severity_level* – Specifies the level of information to be logged from the following list which is ordered from highest to lowest:
  - critical - display critical events
  - error - display error events and all events with a higher severity level
  - warning - display warning events and all events with a higher severity level
  - unusual - display unusual events and all events with a higher severity level
  - info - display info events and all events with a higher severity level
  - trace - display trace events and all events with a higher severity level
  - debug - display all events

Note This keyword is only supported in conjunction with the **active** keyword.

- **critical-info** – Specifies that events with a category attribute of critical information are to be displayed. Examples of these types of events can be seen at bootup when system processes and tasks are being initiated. This is the default setting.

  no-critical-info – Specifies that events with a category attribute of critical information are not to be displayed.

Note These keywords are only supported in conjunction with the **active** keyword.

Important To enable logging of a single instance of a facility, you must first disable all instances of the facility (**logging filter disable facility** *facility* **all**) and then enable logging of the specific instance (**logging filter enable facility** *facility* **instance** *instance_number*. To restore default behavior you must re-enable logging of all instances (**logging filter enable facility** *facility* **all**).

You can display the instance numbers for enabled instances per facility using the Exec mode **show instance-logging** command.

Global Configuration Mode Filtering

You can filter the contents of event logs at the Exec mode and Global Configuration mode levels.

Follow the example below to configure run time event logging parameters for the system:

```
configure
  logging filter runtime facility facility level report_level
  logging display { event-verbosity | pdu-data | pdu-verbosity }
  end
```

Notes:

- **facility** *facility* and **level** *severity_level* – Configure the logging filter that determines which system facilities should be logged and at what levels.
- Repeat for every facility that you would like to log.
- *Optional:* Configure event ID restrictions by adding the **logging disable eventid** command. The system provides the ability to restrict the sending of a specific event ID or a range of event IDs to minimize the amount of data logged to that which is most useful. Repeat to disable logging for additional event IDs or event ID ranges.
- If an administrator restricts event logging for an Event ID or Event ID range using the above command (**logging disable eventid**), the system will generate a Critical Event log "cli 30999 critical" as well as an SNMP trap "1361 (DisabledEventIDs)" with the specific Event IDs or Event ID range that was disabled.

  These event logs and traps are enabled by default in this release, and cannot be disabled.
- If an administrator lowers the logging level (using the **logging filter runtime facility *facility* level *report_level*** command below the default level of "error", the system will generate a Critical Event log "cli 30998 critical" as well as an SNMP trap "1362 (LogLevelChanged)" with the specific Event IDs or Event ID range that was disabled.

  These event logs and traps are enabled by default in this release, and cannot be disabled.

The following examples show the CLI output of the traps generated when event logging or logging levels are changed.

```
show snmp trap statistics
SNMP Notification Statistics:
  ...
Trap Name                           #Gen #Disc  Disable Last Generated
---------------------------------- ----- -----  ------- --------------------
…
DisabledEventIDs                       1     0       0  2021:05:11:15:35:25
LogLevelChanged                        2     0       0  2021:05:11:15:28:03
```

```
show snmp trap history
There are x historical trap records (5000 maximum)

Timestamp                Trap Information
-------------------------
```

```
…
Thu May 11 15:28:03 2021 Internal trap notification 1362 (LogLevelChanged) Logging level
of facility resmgr is changed to critical by user #initial-config# context local privilege
 level Security Administrator ttyname /dev/pts/0 address type IPV4 remote ip address
209.165.202.129
…
Thu May 11 15:35:25 2021 Internal trap notification 1361 (DisabledEventIDs) Event IDs from
 100 to 1000 have been disabled by user adminuser context context privilege level security
 administrator ttyname tty address type IPV4 remote ip address 209.165.202.134
…
Mon May 15 10:14:56 2021 Internal trap notification 1362 (LogLevelChanged) Logging level
of facility sitmain is changed to critical by user staradmin context local privilege level
 Security Administrator ttyname /dev/pts/1 address type IPV4 remote ip address 209.165.202.120
```

Configuring Syslog Servers

Syslog Architecture

System Logging (syslog) is the architecture which produces and sends event information from the MSUP over the UDP transport layer to a centralized Event Message Collector. Syslog uses a client/server architecture:

- **Syslog Client:** A set of processes running on UPF, which operate as the sending device for event messages.
- **Syslog Server:** An external server configured to receive the event messages sent from the MSUP.

MSUP transports event messages using the Syslog Protocol without expecting acknowledgment of receipt. The system forwards event messages regardless if a Syslog Server is available to receive the messages.

Configuring the System to Sent Event Messages to an External Syslog Server

Information that is generated by the run time event logging filters can be transmitted to a syslog server for permanent storage.



Important The data transmitted to the Syslog server is meant to be used for informational purposes. Functions such as billing and performance monitoring should not be based on syslogs.



Important Although the system provides the flexibility to configure syslog servers on a context-by-context basis, it is recommended that all servers be configured in the *local* context in order to isolate the log traffic from the network traffic.

Use the following example to configure syslog servers:

```
configure
  context context_name
logging syslog ip_address [ event-verbosity { min | concise | full } |
facility facilities |  msg-format { rfc3164 | rfc5424 } | pdu-data { none
| hex | hex-ascii } | pdu-verbosity pdu_level |  port number  |  rate value
]
    end
```

NOTES:

- **syslog** *ip_address*: Specifies the IP address of a system log server on the network in IPv4 dotted-decimal or IPv6 colon-separated-hexadecimal notation.
- **event-verbosity { min | concise | full }**: Specifies the level of detail to use in logging of events. Detail level must be one of the following:
  - **min**: Displays minimal detail.
  - **concise**: Displays summary detail.
  - **full**: Displays full detail.
- **facility** *facilities*: Specifies the local facility for which the system logging server's logging options are applied. Local facility must be one of the following:
  - **local0**—Pertains to syslog severity level of 0, Emergency
  - **local1**—Pertains to syslog severity level of 1, Alert
  - **local2**—Pertains to syslog severity level of 2, Critical
  - **local3**—Pertains to syslog severity level of 3, Error
  - **local4**—Pertains to syslog severity level of 4, Warning
  - **local5**—Pertains to syslog severity level of 5, Notice
  - **local6**—Pertains to syslog severity level of 6, Informational
  - **local7**—Pertains to syslog severity level of 7, Debug

  Default: **local7**

  If local facility is not specified, then **local7** is applied by default.

  Multiple system log servers can share the logging options of a given local facility. This allows for the logical grouping of system log servers and the options which affect all of those associated with the same local facility.
- **msg-format { rfc3164 | rfc5424 }**: Configures the message format for each system log server as per RFC3164 or RFC5424. Default: rfc3164.

- **pdu-data { none | hex | hex-ascii }**: Specifies output format for packet data units when logged. Format must be one of the following:
  - **none**: Displays data in raw format.
  - **hex**: Displays data in hexadecimal format.
  - **hex-ascii**: Displays data in hexadecimal and ASCII format (similar to a mainframe dump).
- **pdu-verbosity** *pdu_level*: Specifies the level of verboseness to use in logging of packet data units as a value from 1 through 5, where 5 is the most detailed.
- **port** *number* : Specifies an alternate port number for the system log server. Default: 514.

  number must be an integer value from 1 through 65535.
- **rate** *value*: Specifies the rate at which log entries are allowed to be sent to the system log server. No more than the number specified by *value* is sent to a system log server within any given one-second interval.

  value must be an integer from 0 through 100000. Default: 1000
- Repeat as necessary to configure extra syslog servers. There is no limit to the number of syslog servers that can be configured.

Configuring Active Logs

Active logs are event logs that are operator configurable on a CLI instance-by-CLI instance basis. Active logs configured by an administrative user in one CLI instance are not displayed to an administrative user in a different CLI instance. Each active log can be configured with filter and display properties that are independent of those configured globally for the system. Active logs are displayed in real time as they are generated.

Active logs are not written to the active memory buffer by default. To write active logs to the active memory buffer execute the following command in the Global Configuration mode:

```
logging runtime buffer store all-events
```

When active logs are written to the active memory buffer, they are available to all users in all CLI instances.

Use the following example to configure active logging in Global Configuration mode:

```
logging filter runtime facility facility level report_level
```

NOTES:

- Configure the logging filter that determines which system facilities should be logged and at what levels.
- Repeat for every facility that you want to log.
- *Optional:* Configure event ID restrictions by adding the **logging disable eventid** command. The system provides the ability to restrict the sending of a specific event ID or a range of event IDs to minimize the amount of data logged to that which is most useful. Repeat to disable logging for more event IDs or event ID ranges.

Once all the necessary information has been gathered, the Active log display can be stopped by entering the following command in the EXEC mode:

```
no logging active
```

Specifying Facilities

The following facilities can be configured for logging event data:

- **a10**: A10 interface logging facility
- **a11**: A11 interface logging facility
- **a11mgr**: A11 Manager logging facility
- **aaa-client** : Accounting and authentication client logging facility
- **aaamgr**: Accounting and authentication manager logging facility
- **aaaproxy**: AAAProxy logging facility
- **aal2**: AAL2 protocol logging facility
- **acl-log**: Access Control List logging facility
- **acsctrl**: Active charging Subsystem Controller
- **acsmgr**: Active charging Subsystem Manager
- **alarmctrl**: Alarm Controller logging facility
- **alcap**: ALCAP protocol logging facility
- **alcapmgr**: ALCAPMgr logging facility
- **all**: All facilities
- **apimgr**: API Manager logging facility
- **bfd**: BFD protocol logging facility
- **bgp**: BGP protocol logging facility
- **bindmux**: IPCF BindMux Manager
- **bngmgr**: BNGMGR (BNG DemuxMgr) logging facility
- **bssap+**: Bssap+ logging facility
- **bssgp**: BSSGP Protocol logging facility
- **bulkstat**: Statistics logging facility
- **cactrl**: CACTRL application logging facility
- **callhome**: CALLHOME application logging facility
- **cap**: CAP Protocol logging facility
- **cbsmgr**: CBSMGR logging facility
- **cdf**: CDF logging facility

- **cfctrl**: Content filtering controller
- **cfgmgr**: Config manager logging facility
- **cfmgr**: Content filtering manager
- **cgw**: Converged Gateway facility
- **chmgr**: Checkpoint Manager logging facility
- **cli**: CLI logging logging facility
- **cmp**: CMPv2 logging facility
- **confdmgr**: ConfD Manager - Proclet for communication with confd
- **connectedapps**: CONNECTEDAPPS library logging facility
- **connproxy**: Connection Proxy logging facility
- **credit-control**: Credit Control facility
- **cscf**: IMS/MMD CSCF
- **cscfcpmgr**: CSCFCPMGR logging facility
- **cscfmgr**: SIP CSCF Manager facility
- **cscfnpdb**: CSCF NPDB logging facility
- **cscfrtcp**: MS/MMD CSCF RTCP log facility
- **cscfrtp**: IMS/MMD CSCF RTP log facility
- **cscfttmgr**: SIP CSCFTT Manager facility
- **csp**: Card Slot Port controller logging facility
- **css**: CSS facility
- **css-sig**: CSS RADIUS Signaling facility
- **cx-diameter**: Cx Diameter messages
- **data-mgr**: Data Manager Framework logging facility
- **dcardctrl**: Daughtercard Controller logging facility
- **dcardmgr**: Daughtercard Manager logging facility
- **demuxmgr**: Demux Manager API facility
- **dgmbmgr**: Diameter Gmb Application Manager logging facility
- **dhcp**: Dynamic Host Configuration Protocol logging facility
- **dhcpv6**: DHCPV6
- **dhost**: Distributed Host logging facility
- **diabase**: Diabase messages
- **diactrl**: Diactrl proclet logging

- **diameter**: Diameter endpoint logging facility
- **diameter-acct** : Diameter Accounting
- **diameter-auth**: Diameter Authentication
- **diameter-dns**: Diameter DNS subsystem
- **diameter-ecs**: ECS Diameter signaling
- **diameter-engine** : Diameterv2 engine logging facility
- **diameter-hdd**: Diameter HDD Interface
- **diameter-svc**: Diameter Service
- **diamproxy** : DiamProxy logging facility
- **doulosuemgr**: Doulos UE Manager facility
- **dpath**: IPSEC Data Path logging facility
- **drvctrl**: Driver Controller logging facility
- **eap-ipsec** : EAP
- **eap-sta-s6a-s13-s6b-diameter**: EAP/STA/S6A/S13/S6B Diameter messages
- **ecs-css**: ACSMGR <-> Session Manager Signalling Interface Logging facility
- **egtpc**: EGTP Control plane
- **egtpmgr**: EGTP Demux manager
- **egtpu**: EGTP User plane
- **embms**: eMBMS Gateway facility
- **emctrl**: Element Manager Controller logging facility
- **epdg**: EPDG logging facility
- **event-notif** : Event Notification Interface logging facility
- **evlog**: Event-logger logging facility
- **famgr**: Foreign Agent manager logging facility
- **fapi**: VPP Fastpath API logging facility
- **femgr**: Forwarding Element Manager logging facility
- **firewall**: Firewall logging facility
- **fng**: FNG logging facility
- **gbmgr**: SGSN Gb Interface Manager facility
- **gmm**: GMM Protocol logging facility
- **gprs-app**: GPRS Application logging facility
- **gprs-ns**: GPRS-NS Protocol logging facility

- **gq-rx-tx-diameter** : Gq/Rx/Tx Diameter messages
- **gss-gcdr**: GSS GCDR facility
- **gtpc**: GTP-C Protocol logging facility
- **gtpcmgr**: GTP-C Protocol Manager logging facility
- **gtpp**: GTP-PRIME Protocol logging facility
- **gtpu**: GTP-U Protocol logging facility
- **gtpumgr**: GTPU Demux manager
- **gx-ty-diameter**: Gx/Ty Diameter messages
- **gy-diameter**: Gy Diameter messages
- **hamgr**: Home Agent manager logging facility
- **hat**: High Availability Task process logging facility
- **hdctrl**: HD Controller logging facility
- **henbapp**: HENB App facility
- **henbgw** : HENBGW facility
- **henbgw-pws**: Henbgw Public Warning System Data logging facility
- **henbgw-sctp-acs**: HENBGW access sctp facility
- **henbgw-sctp-nw**: HENBGW network sctp facility
- **henbgwdemux**: HENBGW Demux facility
- **henbgwmgr**: ENBGW Manager facility
- **hnb-gw** : HNB-GW (3G Femto GW) logging facility
- **hnbmgr**: HNBMGR (3G Femto GW DemuxMgr) logging facility
- **hss-peer-service**: HSS Peer Service facility
- **iftask**: iftask logging facility
- **igmp**: IGMP
- **ikev2**: IKEv2
- **ims-authorizatn**: IMS Authorization Service facility
- **ims-sh**: HSS SH Service facility
- **imsimgr** : SGSN IMSI Manager facility
- **imsue**: IMSUE
- **ip-arp** : IP Address Resolution Protocol logging facility
- **ip-interface**: IP interface logging facility
- **ip-route** : IP route logging facility

- **ipms** :IPMS logging facility
- **ipne**: IPNE facility
- **ipsec**: IP Security logging facility
- **ipsecdemux**: IPsecdemux logging facility
- **ipsg**: IP Service Gateway interface logging facility
- **ipsgmgr**: IP Service Gateway facility
- **kvstore**: KV Store facility
- **l2tp-control**: L2TP Protocol Control logging facility
- **l2tp-data**: L2TP Protocol Data logging facility
- **l2tpdemux**: L2TP Demux Manager logging facility
- **l2tpmgr**: L2TP Protocol Manager logging facility
- **lagmgr**: LAGMGR logging facility
- **lcs**: LCS Service logging facility
- **ldap**: LDAP Messages
- **li**: Lawful intercept logging facility
- **linkmgr**: SGSN/SS7 Links Manager facility
- **llc**: LLC Protocol logging facility
- **lmisf**: Lmisf logging facility
- **local-policy**: Local Policy Service facility
- **location-service** : Location Service facility
- **m3ua** : M3UA Protocol logging facility
- **magmgr**: Mobile Access Gateway logging facility
- **map**: MAP Protocol logging facility
- **mcctrl**: SDN Mobility Controller - Controller logging facility
- **mcdemux**: SDN Mobility Controller - DEMUX logging facility
- **mcmgr**: SDN Mobility Controller - Manager logging facility
- **megadiammgr**: Megadiameter Manager (SLF Service)
- **mme-app**: MME application facility
- **mme-misc** : MME miscellaneous logging facility
- **mmedemux**: MME Demux Manager facility
- **mmemgr**: MME Manager facility
- **mmgr**: Primary Manager logging facility

- **mobile-ip**: Mobile IP control logging facility
- **mobile-ip-data**: Mobile IP data logging facility
- **mobile-ipv6**: Mobile IPv6 control logging facility
- **mpls**: MPLS protocol logging facility
- **mrme**: MRME logging facility
- **mseg-app**: MSEG application logging facility
- **mseg-gtpc**: MSEG GTPC application logging facility
- **mseg-gtpu**: MSEG GTPU application logging facility
- **msegmgr**: MSEGMGR (MSEG DemuxMgr) logging facility
- **msgproxy**: Messege Proxy logging facility
- **mtp2** : MTP2 Service logging facility
- **mtp3**: MTP3 Protocol logging facility
- **multicast-proxy**: Multicast Proxy logging facility
- **nas**: NAS 4G Protocol logging facility
- **npuctrl**: Network Processor Unit Control logging facility
- **npumgr**: Network Processor Unit Manager logging facility
- **npumgr-acl**: NPUMGR ACL logging facility
- **npumgr-drv**: NPUMGR DRV logging facility
- **npumgr-flow**: NPUMGR FLOW logging facility
- **npumgr-fwd**: NPUMGR FWD logging facility
- **npumgr-init** : NPUMGR INIT logging facility
- **npumgr-lc**: NPUMGR LC logging facility
- **npumgr-port**: NPUMGR PORT logging facility
- **npumgr-recovery**: NPUMGR RECOVERY logging facility
- **npumgr-rri** : NPUMGR RRI logging facility
- **npumgr-vpn**: NPUMGR VPN logging facility
- **npusim**: NPUSIM logging facility
- **nrfmgr**: NRF Manager logging facility
- **nsh**: NSH logging facility
- **ocsp**: OCSP logging facility
- ospf: OSPF protocol logging facility
- **ospfv3**: OSPFv3 protocol logging facility

- **p2p**: Peer-to-Peer detection logging facility
- **pagingmgr**: PAGINGMGR logging facility
- **pccmgr**: IPCF PCC Manager library
- **pdg**: PDG logging facility
- **pdif**: PDIF logging facility
- **pgw**: PDN Gateway facility
- **plugin**: Plugin logging facility
- **pmm-app** : PMM application logging facility
- **ppp**: PPP link and packet logging facility
- **pppoe**: PPPOE logging facility
- **proclet-map-frwk**: Proclet mapping frameowrk logging facility
- **push**: VPNMGR CDR push logging facility
- **radius-acct**: RADIUS accounting logging facility
- **radius-auth**: RADIUS authentication logging facility
- **radius-coa**: RADIUS change of authorization and radius disconnect logging facility
- **ranap**: RANAP Protocol logging facility
- **rct**: Recovery Control Task logging facility
- **rdt**: Redirect Task logging facility
- **resmgr**: Resource Manager logging facility
- **rf-diameter**: Rf Diameter messages
- **rip**: RIP protocol logging facility
- **rlf**: RLF logging facility
- **rohc**: RObust Header Compression
- **rsvp**: Reservation Protocol logging facility
- **rua**: RUA (3G Femto GW - RUA messages) logging facility
- **s102**: S102 protocol logging facility
- **s102mgr**: S102Mgr protocol logging facility
- **s1ap**: S1AP Protocol logging facility
- **sabp**: SABP logging facility
- **saegw**: SAE Gateway facility
- **sbc**: SBc protocol logging facility
- **sccp**: SCCP Protocol logging facility

- **sct**: Shared Configuration Task logging facility
- **sctp**: SCTP Protocol logging facility
- **sef_ecs**: sef apis printing facility
- **sess-gr**: SM GR facility
- **sessctrl**: Session Controller logging facility
- **sessmgr**: Session Manager logging facility
- **sesstrc**: Session trace logging facility
- **sfcctrl**: Service Funcion Chain Controller logging facility
- **sfcmgr**: Service Funcion Chain Manager logging facility
- **sft**: Switch Fabric Task logging facility
- **sgmbmgr**: eMBMS Gateway Demux facility
- **sgs**: SGs protocol logging facility
- **sgsn-app**: SGSN-APP interface logging facility
- **sgsn-failures**: SGSN call failures logging facility
- **sgsn-gtpc** : SGSNs GTP-C Protocol logging facility
- **sgsn-gtpu** : SGSNs GTP-U Protocol logging facility
- **sgsn-mbms-bearer** : SGSNs MBMS bearer app (SMGR) logging facility
- **sgsn-misc**: SGSN miscellaneous logging facility
- **sgsn-system** : SGSNs System Components logging facility
- **sgsn-test**: SGSN Tests logging facility
- **sgtpcmgr**: SGSN GTPC Manager facility
- **sgw**: Serving Gateway facility
- **sh-diameter** : Sh Diameter messages
- **sitmain**: System Initialization Task main logging facility
- **slmgr**: Smart License Manager logging facility
- **sls**: SLs protocol logging facility
- **sm-app**: SM Protocol logging facility
- **sms**: SMS Service logging facility
- **sndcp**: SNDCP Protocol logging facility
- **snmp**: SNMP logging facility
- **sprmgr**: IPCF SPR Manager library
- **srdb**: Static Rating Database

- **srp**: Service Redundancy Protocol logging facility
- **sscfnni**: SSCFNNI Protocol logging facility
- **sscop**: SSCOP Protocol logging facility
- **ssh-ipsec**: SSH IP Security logging facility
- **ssl**: SSL (Secure socket layer messages) logging facility
- **stat**: Statistics logging facility
- **supserv**: Suplementary Services logging facility
- **sx**: SX Manager logging facility
- **sxdemux**: SX Demux Manager logging facility
- **system**: System logging facility
- **tacacsplus**: TACACS+ Protocol logging facility
- **tcap**: TCAP Protocol logging facility
- **tcpdemux**: TCP Demux Manager logging facility
- **testctrl**: Test Controller logging facility
- **testmgr**: Test Manager logging facility
- **threshold**: Threshold logging facility
- **tsmon**: Traffic Steering Monitor Server facility
- **ttg**: TTG logging facility
- **tucl**: TUCL logging facility
- **udr**: UDR facility
- **uplane**: User Plane logging facility
- **user-data** : User Data logging facility
- **user-l3tunnel**: User layer-3 tunnel logging facility
- **usertcp-stack**: User Tcp Stack
- **vnfma** : VNFM Proxy Agent logging facility
- **vpn**: Virtual Private Network logging facility
- **vpp** : vpp logging facility
- **wimax-data** : WiMAX DATA
- **wimax-r6**: WiMAX R6
- **wsg**: WSG logging facility
- **x2gw-app** : X2GW application logging facility
- **x2gw-demux**: X2GW demux task logging facility

Configuring Trace Logging

Trace logging is useful for quickly resolving issues for specific sessions that are currently active. They are temporary filters that are generated based on a qualifier that is independent of the global event log filter configured using the **logging filter** command in the Exec mode. Like event logs, however, the information generated by the logs is stored in the active memory buffer.

All debug level events associated with the selected call are stored.



Important Trace logs impact session processing. They should be implemented for debug purposes only.

Use the following example to configure trace logs in the Exec mode:

logging trace { callid *call_id* **| ipaddr** *ip_address* **| msid** *ms_id* **| username** *username* **}**

Once all of the necessary information has been gathered, the trace log can be deleted by entering the following command:

no logging trace { callid *call_id* **| ipaddr** *ip_address* **| msid** *ms_id* **| username** *username* **}**

Configuring Monitor Logs

Monitor logging records all activity associated with all of a particular subscriber's sessions. This functionality is available in compliance with law enforcement agency requirements for monitoring capabilities of particular subscribers.

Monitors can be performed based on a subscriber's MSID or username, and are only intended to be used for finite periods of time as dictated by the law enforcement agency. Therefore, they should be terminated immediately after the required monitoring period.

This section provides instructions for enabling and disabling monitor logs.

Enabling Monitor Logs

Use the following example to configure monitor log targets:

```
configure
  logging monitor { ip_addr | ipv6_addr | msid id | username name }
  end
```

Repeat to configure additional monitor log targets.

Disabling Monitor Logs

Use the following example to disable monitor logs:

```
configure
  no logging monitor { ip_addr | ipv6_addr | msid id | username name }
  end
```

Viewing Logging Configuration and Statistics

Logging configuration and statistics can be verified by entering the following command from the Exec mode:

show logging [active | verbose]

When no keyword is specified, the global filter configuration is displayed as well as information about any other type of logging that is enabled.

The following table provides information and descriptions of the statistics that are displayed when the **verbose** keyword is used.

Table 113: Logging Configuration and Statistics Commands

Field	Description
General Logging Statistics	
Total events received	Displays the total number of events generated by the system.
Number of applications receiving events	Displays the number of applications receiving the events.
Logging Source Statistics	
Event sequence ids by process	Displays a list of system processes that have generated events and the reference identification number of the event that was generated.
Msg backlog stat with total cnt	Displays the number of event messages that have been back logged in comparison to the total number of events generated.
LS L2 filter drop rate	Displays the percentage of logging source (LS) layer 2 (L2) event drops.
Abnormal Log Source Statistics	Displays abnormal logging source (LS) statistics, if any.
Runtime Logging Buffer Statistics	
Active buffer	Displays the number of events currently logged in the active memory buffer and a timestamp for the oldest and most recent entries in the buffer.
Inactive buffer	Displays the number of events currently logged in the inactive memory buffer.

Viewing Event Logs Using the CLI

Event logs generated by the system can be viewed in one of the following ways:

- **From the syslog server:** If the system is configured to send logs to a syslog server, the logs can be viewed directly on the syslog server.
- **From the system CLI:** Logs stored in the system memory buffers can be viewed directly from the CLI.

- **From the console port:** By default, the system automatically displays events over the console interface to a terminal provided that there is no CLI session active.

This section provides instructions for viewing event logs using the CLI. These instructions assume that you are at the root prompt for the Exec mode.

Procedure

Step 1 Copy the active log memory buffer to the inactive log memory buffer.

When the active log memory buffer is copied to the inactive log memory buffer existing information in the inactive log memory buffer is deleted.

Both active and inactive event log memory buffers can be viewed using the CLI in Exec mode. However, it is preferable to view the inactive log in order to prevent any data from being over-written. The information from the active log buffer can be copied to the inactive log buffer by entering the following command:

```
logs checkpoint
```

Step 2 View the logs by entering the following command:

```
show logs
```

Configuring and Viewing Crash Logs



Note In the CN-5G 2025.02 release, crash enable URL is not configured, crash cores will be placed at `/var/lib/systemd/coredump/`

In the unlikely even of a software crash, the system stores information that could be useful in determining the reason for the crash. This information can be maintained in system memory or it can be transferred and stored on a network server.

The system supports the generation of the following two types of logs:

- **Crash log:** Crash logs record all possible information pertaining to a software crash (full core dump). Due to their size, they can not be stored in system memory. Therefore, these logs are only generated if the system is configured with a Universal Resource Locator (URL) pointing to a local device or a network server where the log can be stored.
- **Abridged crash log:** Crash event records are automatically generated when a software crash occurs and are stored in flash memory on management cards. The abridged crash log contains a list crash event records along with associated dump files. This log allows you to view event records and dump files via CLI commands.

Crash Logging Architecture

The crash log is a persistent repository of crash event information. Each event is numbered and contains text associated with a CPU (minicore), NPU, or kernel crash. The logged events are recorded into fixed-length records and stored in /flash/crashlog2.

Whenever a crash occurs, the following crash information is stored:

1. The event record is stored in /flash/crashlog2 file (the crash log).
2. The associated minicore, NPU, or kernel dump file is stored in the /flash/crsh2 directory.
3. A full core dump is stored in a user-configured directory.

Configuring Software Crash Log Destinations

The system can be configured to store software crash log information to any of the following locations:

- On VPC
  - **Flash memory:** Accessible by the virtual machine
  - **USB memory stick:** Installed in the USB slot of the platform (USB slot has been enabled via the hypervisor)
- **Network Server:** Any workstation or server on the network that the system can access using the Trivial File Transfer Protocol (TFTP), the File Transfer Protocol (FTP), the Secure File Transfer Protocol (SFTP), or the Hyper-Text Transfer Protocol (HTTP); this is recommended for large network deployments in which multiple systems require the same configuration

Crash log files (full core dumps) are written with unique names as they occur to the specified location. The name format is *crash-card-cpu-time-core*. Where *card* is the card slot, *cpu* is the number of the CPU on the card, and *time* is the Portable Operating System Interface (POSIX) timestamp in hexadecimal notation.

Use the following example to configure a software crash log destination in the Global Configuration mode:

```
configure
  crash enable [ encrypted ] url crash_url
  end
```

NOTES:

- Repeat to configure additional software crash log destinations. There is no limit to the number of destinations that can be configured.

Viewing Abridged Crash Log Information Using the CLI

You can view abridged crash information that is stored as a set of event records in flash memory on management cards (**/flash/crashlog2**). Each crash event record has an associated dump file (minicore, NPU or kernel) that can also be displayed (**/flash/crsh2**)

Follow the instructions in this section to view software crash events that have occurred on the system. These instructions assume that you are at the root prompt for the Exec mode.

Procedure

Step 1 View a list of software crash events by entering the following Exec mode command:

show crash { all | list | number *crash_num* **}**

NOTES:

- Run **show crash list** to obtain the number for a specific crash event.
- Run **show crash number** *crash_num* to display the output for the target crash event.

The resulting output may not be the same for all platforms:

Information about similar crash events is suppressed in the output of this command.

Step 2 View the dump file associated with a specific crash event.

The information contained in the dump file helps identify and diagnose any internal or external factors causing the software to crash.

- Crash # – unique number assigned by StarOS when logging the crash event
- SW Version – StarOS build release in format: RR.n(bbbbb)
- Similar Crash Count – number of similar crashes
- Time of first crash – timestamp when first crash occurred in format: YYYY-MMM-DD+hh:mm:ss
- Failure message – text of event message
- Function – code identifier
- Process – where the crash occurred (Card, CPU, PID, etc.)
- Crash time – timestamp for when the crash occurred in the format: YYYY-MMM-DD+hh:mm:ss time zone
- Recent errno – text of most recent error number.
- Stack – memory stack information
- Last Bounce – information about the messaging received prior to the crash
- Registers – memory register contents
- Current inbound message – hexadecimal information for the current inbound message
- Address Map
- Recent heap activity (oldest first)
- Recent events (oldest first)
- Profile depth

Reducing Excessive Event Logging

Event logging (evlogd) is a shared medium that captures event messages sent by StarOS facilities. When one or more facilities continuously and overwhelmingly keep sending a high volume of event messages, the

remaining non-offender facilities are impacted. This scenario degrades system performance, especially as the number of facilities generating logs increases.

Rate-control of event message logging is handled in the Log Source path. Essentially, every second a counter is set to zero and is incremented for each log event that is sent to evlogd. If the count reaches a threshold before the second is up, the event is sent, queued, or dropped (if the evlogd messenger queue is full).

When any facility exceeds the upper threshold set with this command for the rate of message logging and remains in the same state for prolonged interval, StarOS notifies the user through an SNMP trap or alarm.

A new threshold command allows you to specify the percentage of facility event queue full. When this threshold is exceeded, an SNMP trap and alarm are generated that specifies the offending facility.

The formats for the SNMP traps that are associated with this command are as follows:

- **ThreshLSLogsVolume**

```
<timestamp> Internal trap notification <trap_id> (ThreshLSLogsVolume) threshold
<upper_percent>%
measured value <actual_percent>% for facility <facility_name> instance <instance_id>
```

- **ThreshClearLSLogsVolume**

```
<timestamp> Internal trap notification <trap_id> (ThreshClearLSLogsVolume) threshold
<upper_percent>%
measured value <actual_percent>% for facility <facility_name> instance <instance_id>
```

If a trigger condition occurs within the polling interval, the alert or alarm is not generated until the end of the polling interval.

Both traps can be enabled or suppressed through the Global Configuration mode **snmp trap** command.

Configuring Log Source Thresholds

There are three Global Configuration mode commands associated with configuring and implementing Log Source thresholds.

1. **threshold ls-logs-volume** – sets the parameters for the upper and lower thresholds for generating and clearing traps/alarms respectively.
2. **threshold poll ls-logs-volume interval** – establishes the polling interval for this threshold.
3. **threshold monitoring ls-logs-volume** – turns monitoring of this threshold on and off.

Use the following example to configure syslog servers:

```
configure
  [ default ] threshold ls-logs-volume upper_percent [ clear lower_percent ]
  [ default ] threshold poll ls-logs-volume interval duration
  [ no ] threshold monitoring ls-logs-volume
  end
```

Notes:

- *upper_percent* and *lower_percent* are expressed as integers from 0 to 100. Default value for *upper_percent* is 90%. If *lower_percent* is not specified, the default clear value is *upper_percent*.
- **threshold poll ls-logs-volume interval** sets the polling interval in seconds. The default interval is 300 seconds (5 minutes).

- **threshold monitoring ls-logs-volume** enables or disables this feature.

You can verify the configuration of this threshold by running the Exec mode **show threshold** command.

Checkpointing Logs

Checkpointing identifies logged data as previously viewed or marked. Checkpointing allows you to only display log information since the last checkpoint.

Individual logs may have up to 50,000 events in the active log. Checkpointing the logs results in at most 50,000 events being in the inactive log files. This gives a maximum of 100,000 events in total which are available for each facility logged.

Checkpoint log data through the EXEC mode logs checkpoint command to set the log contents to a well-known point prior to special activities taking place. This command may also be a part of periodic regular maintenance to manage log data.

Checkpointing logs moves the current log data to the inactive logs. Only the most recently check pointed data is retained in the inactive logs. A subsequent check pointing of the logs results in the prior check pointed inactive log data being cleared and replaced with the newly check pointed data. Checkpointed log data is not available for viewing.



Important Checkpointing logs should be done periodically to prevent the log files becoming full. Logs which have 50,000 events logged discard the oldest events first as new events are logged.

Saving Log Files

Log files can be saved to a file in a local or remote location specified by a URL. Use the following EXEC mode command to save log files:

save logs { *url* } **[active]] [inactive] [callid** *call_id* **]**
[event-verbosity *evt_verboseness* **] [facility** *facility* **] [level** *severity_level* **]**
[pdu-data *pdu_format* **] [pdu-verbosity** *pdu_verboseness* **] [since** *from_date_time*
[until *to_date_time* **]] [| { grep** *grep_options* **| more }]**

OPTIONS:

- *url*: Specifies the location to store the log file(s). *url* may refer to a local or a remote file and must be entered in the following format.
- **active**: Saves data from active logs.
- **inactive**: Saves data from inactive logs.
- **callid** *call_id*: Specifies a call ID for which log information is to be saved as a 4-byte hexadecimal number.
- **event-verbosity** *evt_verboseness*: Specifies the level of verboseness to use in the displaying of event data as one of:

  - *min*: Logs minimal information about the event. Information includes event name, facility, event ID, severity level, date, and time.
  - *concise*: Logs detailed information about the event, but does not provide the event source within the system.
  - *full*: Logs detailed information about event, including source information, identifying where within the system the event was generated.

- **facility** *facility*: Specifies the facility to modify the filtering of logged information.
- **level** *severity_level*: Specifies the level of information to be logged in the following list which is ordered from highest to lowest:
  - *critical*: Logs critical events
  - *error*: Logs error events and all events with a higher severity level
  - *warning*: Logs warning events and all events with a higher severity level
  - *unusual*: Logs unusual events and all events with a higher severity level
  - *info*: Logs info events and all events with a higher severity level
  - *trace*: Logs trace events and all events with a higher severity level
  - *debug*: Logs all events
- **pdu-data** *pdu_format*: Specifies output format for the display of packet data units as one of:
  - *none* - raw format (unformatted).
  - *hex* - hexadecimal format.
  - *hex-ascii* - hexadecimal and ASCII similar to a main-frame dump.
- **pdu-verbosity** *pdu_verboseness* : Specifies the level of verboseness to use in the displaying of packet data units as a value from 1 to 5, where 5 is the most detailed.
- **since** *from_date_time*: Saves only the log information which has been collected more recently than *from_date_time*
- **until** *to_date_time*: Saves no log information more recent than *to_date_time*. Defaults to current time when omitted.
- *from_date_time* and *to_date_time* must be formatted as YYYY:MM:DD:HH:mm or YYYY:MM:DD:HH:mm:ss. Where:
  - YYYY = 4-digit year
  - MM = 2-digit month in the range 01 through 12
  - DD = 2-digit day in the range 01 through 31
  - HH = 2-digit hour in the range 00 through 23
  - mm = 2-digit minute in the range 00 through 59
  - ss = 2 digit second in the range 00 through 59

to_date_time must be a time which is more recent than *from_date_time*.

Using the **until** keyword allows for a time range of log information; using only the **since** keyword will display all information up to the current time.

- **grep** *grep_options* | **more**: Pipes (sends) the output of this command to the specified command. You must specify a command to which the output of this command is sent.

Event ID Overview

Important The use of event IDs depends on the platform type and the licenses running on the platform.

Identification numbers (IDs) are used to reference events as they occur when logging is enabled on the system. As described previously, logs are collected on a per facility basis. Each facility possesses its own range of event IDs as indicated in the following table.

Table 114: System Facilities and Event ID Ranges

Facility	Description	Event ID Range
a10	A10 Protocol Facility	28000-28999
a11	A11 Protocol Facility	29000-29999
a11mgr	A11 Manager Facility	9000-9999
aaa-client	AAA Client Facility	6000-6999
aaamgr	AAA Manager Facility	36000-36999
aaaproxy	AAA Proxy Facility	64000-64999
aal2	AAL2 Protocol Facility	173200-173299
acl-log	IP Access Control List (ACL) Facility	21000-21999
acsctrl	Active Charging Service Controller (ACSCtrl) Facility	90000-90999
acsmgr	Active Charging Service Manager (ACSMgr) Facility	91000-91999
afctrl	Ares Fabric Controller (ASR 5500 only)	186000-186999
afmgr	Ares Fabric Manager (ASR 5500 only)	187000-187999
alarmctrl	Alarm Controller Facility	65000-65999
alcap	Access Link Control Application Part (ALCAP) Protocol Facility	160900-161399
alcapmgr	ALCAP Manager Facility	160500-160899
asf	ASF Facility	73000-73999
asfprt	ASFPRT Facility	59000-59999
asngwmgr	Access Service Network (ASN) Gateway Manager Facility	100000-100499
asnpcmgr	ASN Paging/Location-Registry Manager Facility	100500-100999

Facility	Description	Event ID Range
bcmcs	Broadcast/Multicast Service (BCMCS) Facility	109000-109999
bfd	Bidirectional Forwarding Detection (BFD) Protocol Facility	170500-170999
bgp	Border Gateway Protocol (BGP) Facility	85000-85999
bindmux	BindMux Manager Facility [Intelligent Policy Control Function (IPCF)]	158200-158999
bngmgr	Broadband Network Gateway (BNG) Manager Facility	182000-182999
bssap	Base Station System Application Part+ (BSSAP+) Service Facilities	131000-131199
bssgp	Base Station System GPRS Protocol (BSSGP) Facility	115050-115099
callhome	Call Home Facility	173600-173999
cap	CAMEL Application Part (CAP) Facility	87900-88099
chatconf	CHATCONF Facility	74000-74999
cli	Command Line Interface (CLI) Facility	30000-30999
connproxy	Connection Proxy Facility	190000-190999
crdt-ctl	Credit Control Facility	127000-127999
csg	Closed Subscriber Groups (CSG) Facility	188000-188999
csg-acl	CSG Access Control List (ACL) Facility	189000-189999
csp	Card/Slot/Port (CSP) Facility	7000-7999
css	Content Steering Service (CSS) Facility [ESC]	77000-77499
css-sig	Content Service Selection (CSS) RADIUS Signaling Facility	77500-77599
cx-diameter	Cx Diameter Message Facility	92840-92849
dcardctrl	Daughter Card Controller Facility	62000-62999
dcardmgr	Daughter Card Manager Facility	57000-57999
demuxmgr	Demux Manager Facility	110000-110999
dgmbmgr	Diameter Gmb (DGMB) Application Manager Facility	126000-126999
dhcp	DHCP Facility	53000-53999
dhcpv6	DHCPv6 Protocol Facility	123000-123999
dhost	Distributed Host Manager Facility	83000-83999
diameter	Diameter Endpoint Facility	92000-92599
diabase	Diabase Message Facility	92800-92809
diameter-acct	Diameter Accounting Protocol Facility	112000-112999
diameter-auth	Diameter Authentication Protocol Facility	111000-111999
diameter-dns	Diameter DNS Subsystem Facility	92600-92699
diameter-ecs	ECS Diameter Signaling Facility	81990-81999

Facility	Description	Event ID Range
diameter-hdd	Diameter Horizontal Directional Drilling (HDD) Interface Facility	92700-92799
diameter-svc	Diameter Service Facility	121200-121999
diamproxy	Diameter Proxy Facility	119000-119999
dpath	Data Path for IPSec Facility	54000-54999
drvctrl	Driver Controller Facility	39000-39999
ds3mgr	DS3 and DS3/E Line Card Manager Facility (part of NPU Manager Controller Facility)	40000-40999
eap-diameter	Extensible Authentication Protocol (EAP) Diameter Facility	92870-92879
eap-ipsec	EAP IPSec Facility	118000-118999
ecs-css	ACS Session Manager (ACSMgr) Signalling Interface Facility	97000-97099
edr	Event Data Record (EDR) Facility	80000-80999
egtpc	eGTP-C Facility	141000-141999
egtpmgr	eGTP Manager Facility	143000-143999
egtpu	eGTP-U Facility	142000-142999
epdg	Evolved Packet Data Gateway (ePDG) Facility	178000-178999
evlog	Event Log Facility	2000-2999
famgr	Foreign Agent (FA) Manager Facility	33000-33999
firewall	Firewall Facility	96000-96999
fng	Femto Network Gateway (FNG) Facility	149000-149999
gbrmgr	Gb-Manager Facility	201900-202699
gcdr	GGSN-Charging Data Record (G-CDR) Facility	66000-66999
gmm	GPRS Mobility Management (GMM) Facility	88100-88299
gprs-app	General Packet Radio Service (GPRS) Application Facility	115100-115399
gprs-ns	GPRS-NS Protocol Facility	115000-115049
gq-rx-tx-diameter	Gq/Rx/Tx Diameter Messages Facility	92830-92839
gss-gcdr	GTPP Storage Server GCDR Facility	98000-98099
gtpc	GTPC Protocol Facility	47000-47999
gtpcmgr	GTPC Signaling Demultiplexer Manager Facility	46000-46999
gtpp	GTP-PRIME Protocol Facility	52000-52999
gtpu	GTPU Protocol Facility	45000-45999
gtpumgr	GTPU Manager Facility	157200-157999
gx-ty-diameter	Gx/Ty Diameter Messages Facility	92820-92829
gy-diameter	Gy Diameter Messages Facility	92810-92819

Facility	Description	Event ID Range
h248prt	H.248 Protocol Facility	42000-42999
hamgr	Home Agent (HA) Manager Facility	34000-34999
hat	High Availability Task (HAT) Facility	3000-3999
hdctrl	Hard Disk (HD) Controller Facility	132000-132999
hddshare	HDD Share Facility	184000-184999
henb-gw	Home eNodeB-GW Facility	195000-195999
henbapp	Home eNodeB Application Facility	196000-196999
henbgwdemux	Home eNodeB-GW Demux Facility	194000-194999
henbgwmgr	Home eNodeB-GW Manager Facility	193000, 193999
hnb-gw	Home NodeB (HNB) Gateway Facility	151000-151999
hnbmgr	HNB Manager Facility	158000-158199
hss-peer-service	Home Subscriber Server (HSS) Facility [MME]	138000-138999
igmp	Internet Group Management Protocol (IGMP) Facility	113000-113999
ikev2	IKEv2 Facility	122000-122999
ims-authorizatn	IMS Authorization Service Library Facility	98100-98999
ims-sh	IMS SH Library Facility	124000-124999
imsimgr	International Mobile Subscriber Identity (IMSI) Manager Facility	114000-114999
imsue	IMS User Equipment (IMSUE) Facility	144000-145999
ip-arp	IP Address Resolution Protocol (ARP) Facility	19000-19999
ip-interface	IP Interface Facility	18000-18999
ip-route	IP Route Facility	20000-20999
ipms	Intelligent Packet Monitoring System (IPMS) Facility	134000-134999
ipne	IP Network Enabler (IPNE) Facility	192000-192999
ipsec	IPSec Protocol Facility	55000-56998
ipsg	IP Services Gateway (IPSG) Facility	128000-128999
ipsgmgr	IPSG Manager (IPSGMgr) Facility	99000-99999
ipsp	IP Pool Sharing Protocol (IPSP) Facility	68000-68999
kvstore	Key/Value Store (KVSTORE) Facility	125000-125999
l2tp-control	L2TP Control PDU Protocol Facility	50000-50999
l2tp-data	L2TP Data PDU Protocol Facility	49000-49999
l2tpdemux	L2TP Demux Facility	63000-63999
l2tpmgr	L2TP Manager Facility	48000-48999

Facility	Description	Event ID Range
lagmgr	Link Aggregation Group (LAG) Manager Facility	179000-179999
ldap	Lightweight Directory Access Protocol (LDAP) Request Facility	160000-160499
li	Lawful Intercept (LI) Log Facility	69000-69999
linkmgr	Link Manager Facility	89500-89999
llc	Logical Link-Control (LLC) Layer Facility (GPRS)	115700-115799
local-policy	Local Policy Configuration Facility	161400-162399
m3ap	M3 Application Protocol (M3AP) Facility	211500-211999
m3ua	MTP Level 3 (M3UA) Protocol Facility [SIGTRAN]	87500-87699
magmgr	Mobile Access Gateway (MAG) Manager Facility	137500-137999
map	Mobile Application Part (MAP) Protocol Facility [SS7]	87100-87299
megadiammgr	MegaDiameter Manager Facility	121000-121199
mme-app	Mobility Management Entity (MME) Application Facility	147000-147999
mme-embms	MME evolved Multimedia Broadcast Multicast Service (eMBMS) Facility	212000-212499
mme-misc	MME Miscellaneous Facility	155800-156199
mmedemux	MME Demux Manager Facility	154000-154999
mmemgr	MME Manager Facility	137000-137499
mmgr	Master Manager (MMGR) Facility	86000-86399
mobile-ip	Mobile IP (MIP) Protocol Facility	26000-26999
mobile-ip-data	MIP Tunneled Data Facility	27000-27999
mobile-ipv6	Mobile IPv6 Facility	129000-129999
mpls	Multiprotocol Label Switching (MPLS) Facility	163500-163999
mseg-app	Mobile Services Edge Gateway (MSEG) Application Facility Not supported in this release.	172300-172999
mseg-gtpc	MSEG GTPC Application Facility Not supported in this release.	172000-172199
mseg-gtpu	MSEG GTPU Application Facility Not supported in this release.	172200-172299
msegmgr	MSEG Manager Facility Not supported in this release.	171000-171999
mtp2	Message Transfer Part 2 (MTP2) Service Facility [SS7]	116900-116999
mtp3	Message Transfer Part 3 (MTP3) Service Facility [SS7]	115600-115699
multicast-proxy	Multicast Proxy Facility	94000-94999

Facility	Description	Event ID Range
nas	Network Access Signaling (NAS) Facility	153000-153999
netwstrg	Network Storage Facility	78000-78999
npuctrl	Network Processing Unit (NPU) Control Facility	16000-16999
npudrv	NPU Driver Facility	191000-191999
npumgr	NPU Manager (NPUMGR) Facility	17000-17999
npumgr-acl	NPUMGR ACL Facility	169000-169999
npumgr-drv	NPUMGR Driver Facility	185000-185999
npumgr-flow	NPUMGR Flow Facility	167000-167999
npumgr-fwd	NPUMGR Forwarding Facility	168000-168999
npumgr-init	NPUMGR Initialization Facility	164000-164999
npumgr-lc	NPUMGR LC Facility	180000-180999
npumgr-port	NPUMGR Port Facility	166000-166999
npumgr-recovery	NPUMGR Recovery Facility	165000-165999
npumgr-vpn	NPUMGR VPN Facility	181000-181999
npusim	NPUSIM Facility	176000-176999
ntfy-intf	Event Notification Interface Facility	170000-170499
orbs	Object Request Broker (ORB) System Facility	15000-15999
ospf	Open Shortest Path First (OSPF) Protocol Facility	38000-38999
ospfv3	OSPFv3 Protocol Facility [IPv6]	150000-150999
p2p	Peer-to-Peer (P2P) Facility	146000-146999
pccmgr	Policy Charging and Control (PCC) Manager Facility	159000-159499
pdg	Packet Data Gateway (PDG) Facility	152010-152999
pdgdmgr	PDG TCP Demux Manager (pdgdmgr) Facility (this is a customer-specific facility)	162400-162999
pdif	Packet Data Interworking Function (PDIF) Facility	120000-120999
pgw	Packet Data Network Gateway (PGW) Facility	139000-139999
phs	Personal Handyphone System (PHS) Facility	177000-177999
phs-control	PHS X1/X5 and X2/X6 Interface Facility	136000-136899
phs-data	PHS Data Facility	136900-136979
phs-eapol	PHS EAP over LAN (EAPOL) Facility	136980-136999
phsgwmgr	PHS Gateway Manager Facility	135000-135499
phspcmgr	PHS Paging Controller Manager Facility	135500-135999
pmm-app	Packet Mobility Management (PMM) Application Facility [SGSN]	89200-89499

Facility	Description	Event ID Range
ppp	Point-To-Point Protocol (PPP) Facility	25000-25999
pppoe	Point-to-Point Protocol over Ethernet (PPPoE) Facility	183000-183999
ptt	PTT Facility	76000-76999
push	PUSH (VPNMgr CDR Push) Facility	133000-133999
radius-acct	RADIUS Accounting Protocol Facility	24000-24999
radius-auth	RADIUS Authentication Protocol Facility	23000-23999
radius-coa	RADIUS Change of Authorization (CoA) and Disconnect Facility	70000-70999
ranap	Radio Access Network Application Part (RANAP) Facility	87700-87899
rct	Recovery Control Task (RCT) Facility	13000-13999
rdt	Redirector Task (RDT) Facility	67000-67999
resmgr	Resource Manager (RM) Facility	14000-14999
rf-diameter	Rf Diameter Messages Facility	92860-92869
rip	Routing Information Protocol (RIP) Facility	35000-35999
rohc	Robust Header Compression (ROHC) Protocol Facility	103000-103999
rsvp	RSVP Protocol Facility	93000-93999
rua	RANAP User Adaptation (RUA) Protocol Facility	152000-152009
s1ap	S1 Application Protocol (S1AP) Facility	155200-155799
saegw	System Architecture Evolution Gateway Facility	191000-191999
sccp	Signalling Connection Control Part (SCCP) Protocol Facility [SS7]	86700-86899
sct	Shared Configuration Task (SCT) Facility	32000-32099
sctp	Stream Control Transmission Protocol (SCTP) Protocol Facility	87300-87499
sess-gr	SESS-GR Facility	77600-77999
sessctrl	Session Controller Facility	8000-8999
sessmgr	Session Manager Facility	10000-12999
sesstrc	Session Trace Facility	155000-155199
sft	Switch Fabric Task (SFT) Facility	58000-58999
sgs	SGs Interface Protocol Facility [MME]	173000-173199
sgsn-app	SGSN Application Interface Facility	115900-115999
sgsn-failures	SGSN Call Failures Facility	89100-89199
sgsn-gtpc	SGSN GTP-C Protocol Facility	116000-116599
sgsn-gtpu	SGSN GTP-U Protocol Facility	86900-87099
sgsn-mbms-bearer	SGSN MBMS Bearer Application (SMGR) Facility	116600-116799

Facility	Description	Event ID Range
sgsn-misc	SGSN Miscellaneous Facility	88800-89099
sgsn-system	SGSN System Components Facility	86400-86499
sgsn-test	SGSN Tests Facility	88700-88799
sgsn2	SGSN2 Facility	114000-117999
sgtpcmgr	SGSN GTP-C (SGTPC) Manager Facility	117000-117999
sgw	Serving Gateway (SGW) Facility	140000-140999
sh-diameter	Sh Diameter Messages Facility	92850-92859
sipcdprt	SIPCDPRT Facility	95000-95999
sitmain	System Initiation Task (SIT) Main Facility	4000-4999
sm-app	Short Message Service (SMS) Facility	88300-88499
sms	SMS Service Facility	116800-116899
sndcp	Sub Network Dependent Convergence Protocol (SNDCP) Facility	115800-115899
snmp	Simple Network Management Protocol (SNMP) Facility	22000-22999
sprmgr	Subscriber Policy Register (SPR) Manager Facility	159500-159999
srdb	Static Rating Database Facility	102000-102999
srp	Service Redundancy Protocol (SRP) Facility	84000-84999
sscfnni	SSCFNNI Protocol Facility [ATM]	115500-115599
sscop	SSCOP Protocol Facility [ATM]	115400-115499
ssh-ipsec	SSH IP Security Facility	56999-56999
ssl	SSL Facility (this is a customer-specific facility)	156200-157199
stat	Statistics Facility	31000-31999
system	System Facility	1000-1999
tacacs+	TACACS+ Protocol Facility	37000-37999
taclcp	TACLCP Facility	44000-44999
tcap	Transaction Capabilities Application Part (TCAP) Protocol Logging Facility [SS7]	86500-86699
testctrl	Test Controller Facility	174000-174999
testmgr	Test Manager Facility	175000-175999
threshold	Threshold Facility	61000-61999
ttg	Tunnel Termination Gateway (TTG) Facility	130000-130999
tucl	TCP/UDP Convergence Layer (TUCL) Facility [SS7]	88500-88699
udr	User Data Record (UDR) Facility	79000-79999
user-data	User-Data Facility	51000-51999

Facility	Description	Event ID Range
user-l3tunnel	User L3 Tunnel Facility	75000-75999
usertcp-stack	User TCP Stack Facility	173300-173499
vim	Voice Instant Message (VIM) Facility	60000, 60999
vinfo	VINFO Facility	82000, 82999
vmgctrl	Virtual Media Gateway (VMG) Controller Facility	41000, 41999
vmgctxmgr	VMG Context Manager Facility	43000, 43999
vpn	Virtual Private Network (VPN) Facility	5000-5999
wimax-data	WiMAX DATA Facility	104900-104999
wimax-r6	WiMAX R6 Protocol (Signaling) Facility	104000-104899

Event Severities

The system provides the flexibility to configure the level of information that is displayed when logging is enabled. The following levels are supported:

- **critical:** Logs only those events indicating a serious error has occurred that is causing the system tor a system component to cease functioning. This is the highest severity level.
- **error:** Logs events that indicate an error has occurred that is causing the system or a system component to operate in a degraded state. This level also logs events with a higher severity level.
- **warning:** Logs events that may indicate a potential problem. This level also logs events with a higher severity level.
- **unusual:** Logs events that are very unusual and may need to be investigated. This level also logs events with a higher severity level.
- **info:** Logs informational events and events with a higher severity level.
- **trace:** Logs events useful for tracing and events with a higher severity level.
- **debug:** Logs all events regardless of the severity.

Each of the above levels correspond to the "severity" level of the event ID. Therefore, only those event IDs with a "severity" level equal to the logging level are displayed.

Understanding Event ID Information in Logged Output

This section explains the event information that is displayed when logging is enabled.

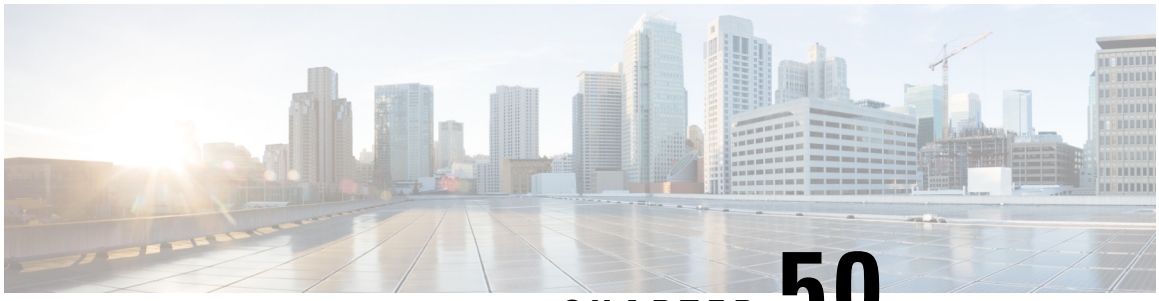
The following displays a sample output for an event that was logged.

```
2011-Dec-11+5:18:41.993 [cli 30005 info] [8/0/609 cli:8000609 _commands_cli.c:1290] [software
 internal system] CLI session ended for Security Administrator admin on device /dev/pts/2
```

The following table describes the elements of contained in the sample output.

Table 115: Event Element Descriptions

Element	Description
2011-Dec-11+5:18:41.993	Date/Timestamp indicating when the event was generated
[cli 30005 info]	Information about the event including: • The facility the event belongs to • The event ID • The event's severity level In this example, the event belongs to the CLI facility, has an ID of 3005, and a severity level of "info".
[8/0/609 cli:8000609 _commands_cli.c:1290]	Information about the specific CLI instance.
[software internal system]	Indicates that the event was generated because of system operation.
CLI session ended for Security Administrator admin on device /dev/pts/2	The event's details. Event details may, or may not include variables that are specific to the occurrence of the event.



CHAPTER 50

UPF Ingress Interfaces

- Feature Summary and Revision History, on page 547
- Feature Description, on page 548
- UPF Ingress Interface Type Configuration, on page 548
- Multiple GTPU IP Addresses Support, on page 554

Feature Summary and Revision History

Summary Data

Table 116: Summary Data

Applicable Product(s) or Functional Area	5G-UPF
Applicable Platform(s)	VPC-SI SMI
Feature Default Setting	Disabled – Configuration Required
Related Changes in this Release	Not Applicable
Related Documentation	*UCC 5G UPF Configuration and Administration Guide*

Revision History

Table 117: Revision History

Revision Details	Release
Added support for binding multiple GTPU IP addresses to ensure even distribution of traffic across the VPP worker threads on UPF.	2023.04.0

Revision Details	Release
Added support for IP separation per interface to allow separate networks for N3 and N9 interfaces.	2023.03.0 2023.02.0
First introduced.	2020.02.0

Feature Description

UPF supports GTP-U ingress interfaces to initiate the user plane service. The UPF supports different GTP-U ingress interfaces to allow a separate network for N3 interface and N9, S5u, and S8u interfaces. The N9, S5u, and S8u interfaces share the same public IP.

The supported GTP-U ingress interfaces include:

- **N3**: N3 is the interface between gNodeB and UPF.
- **N9**: The N9 interface connects two UPFs. It is the interface between intermediate I-UPF and UPF session anchor connecting different PLMN.
- **S5u**: S5u is the interface between SGW-U and PGW-U in the 4G.
- **S8u**: S8u is an inter-PLMN variant of the S5u interface.

The user plane service associates with the GTP-U service using the **associate gtpu-service** CLI command in User Plane Service configuration mode. To enable the **upf-ingress** command, you must configure the **require upf** command on the UPF. You need the UPF license to enable the **require upf** command.

Note UPF does not support multiple GTP-U service selections for a single call.

UPF Ingress Interface Type Configuration

Configuring UPF Ingress Interface Type

To configure the UPF ingress interface type, use the following configuration:

To enable the **upf-ingress** command, you must configure the **require upf** command on the UPF. You need the UPF license to enable the **require upf** command.

```
configure
  context context_name
    user-plane-service service_name
      [ no ] associate gtpu-service gtpu_service_name upf-ingress
interface-type [ n9 | s5u | s8u | n3 | n9-s5u-s8u ]
      end
```

NOTES:

- **associate gtpu-service** *gtpu_service_name*: Associate the GTP-U service with the user plane service.
- **upf-ingress**: Configure the interface type as UPF ingress.
- **interface-type [n9 | s5u | s8u | n3 | n9-s5u-s8u]**: Configure the desired GTP-U ingress interface type.
- If the N3 GTP-U ingress service is not present, UPF uses the default interface.
- If the configuration for N9, S5u, or S8u interfaces is not present, UPF performs the following actions:
  - UPF looks to use the N9 interface first followed by the S5u and S8u interfaces.
  - If there is no GTP-U service, UPF uses the default interface.

Configuration Example

The following is an example configuration:

```
associate gtpu-service n3-ingress upf-ingress
associate gtpu-service n3-ingress upf-ingress interface-type n3
associate gtpu-service n9-ingress upf-ingress interface-type n9
associate gtpu-service s5-ingress upf-ingress interface-type s5
associate gtpu-service s8-ingress upf-ingress interface-type s8
```

Configuring Echo Response for N9 Interface

The echo response in a GTP-U service provides the availability of peer UPF.

The following is a sample echo response configuration for the N9 interface in hUPF:

```
gtpu-service n9-ingress
   echo-interval 60
   echo-retransmission-timeout 1
   bind ipv4-address 60.60.41.5 bearer-type all
   exit
```

The following is a sample echo response configuration for the N9 interface in vUPF:

```
gtpu-service n9-egress
   echo-interval 60
   echo-retransmission-timeout 1
   bind ipv4-address 60.60.42.5 bearer-type all
   exit
```

Separating N3 and N9 BGP Routes

The following is a sample configuration to separate the N3 and N9 BGP routes for hUPF.

```
 config
  context SAEGW
    ip vrf mpls-vrf-1
    #exit
    ip vrf mpls-vrf-2
    #exit
    mpls bgp forwarding
    bfd-protocol
    #exit
    ip prefix-list name up seq 15 permit 50.50.17.4/32
    ip prefix-list name up seq 20 permit 50.50.17.5/32
```

```
ip prefix-list name up seq 25 permit 50.50.17.6/32
ip prefix-list name up seq 40 permit 50.50.17.7/32
ip prefix-list name n9-up1 seq 50 permit 60.60.17.5/32
ip prefix-list name n9-up1 seq 55 permit 60.60.17.6/32
route-map up_routes1 permit 10
  match ip address prefix-list n9-up1
#exit
route-map up_routes permit 10
  match ip address prefix-list up
#exit
router bgp 61435
  maximum-paths ebgp 4
  neighbor 20.14.35.100 remote-as 65000
  neighbor 20.14.35.100 ebgp-multihop
  neighbor 20.14.35.100 update-source 20.14.35.98
  neighbor 2001:4888:121:1435::1 remote-as 65000
  neighbor 2001:4888:121:1435::1 ebgp-multihop
  neighbor 2001:4888:121:1435::1 update-source 2001:4888:121:1435::2
  neighbor 20.14.40.100 remote-as 65000
  neighbor 20.14.40.100 ebgp-multihop
  neighbor 20.14.40.100 update-source 20.14.40.98
  address-family ipv4
    neighbor 20.14.35.100 max-prefix 8192 threshold 50
    neighbor 20.14.35.100 route-map up_routes out
    neighbor 20.14.35.100 send-community extended
    neighbor 20.14.40.100 max-prefix 8192
    neighbor 20.14.40.100 route-map up_routes1 out
    neighbor 20.14.40.100 send-community extended
    redistribute connected
    redistribute static
  #exit
  address-family ipv6
    neighbor 2001:4888:121:1435::1 activate
    redistribute connected
    redistribute static
  #exit
#exit
interface SxC loopback
  ip address 50.50.17.4 255.255.255.255
#exit
interface SxC-ipv6 loopback
  ipv6 address 2001:4888:50:50::17/128
#exit
interface gn-1
  ip address 20.14.35.98 255.255.255.0
  ipv6 address 2001:4888:121:1435::2/64 secondary
#exit
interface gn-2
  ip address 20.14.36.98 255.255.255.0
  ipv6 address 2001:4888:121:1436::2/64 secondary
#exit
interface gn-n9-ingress
  ip address 20.14.40.98 255.255.255.0
#exit
interface n9-ingress-hupf loopback
  ip address 60.60.17.5 255.255.255.255
#exit
interface pgw-gtpu loopback
  ip address 50.50.17.5 255.255.255.255
#exit
interface sgw-gtpu-egress loopback
  ip address 50.50.17.6 255.255.255.255
#exit
interface sgw-gtpu-ingress-phazr loopback
```

```
    ip address 60.60.17.6 255.255.255.255
  #exit
  subscriber default
  exit
  aaa group default
  #exit
  gtpp group default
    gtpp limit-secondary-rat-usage 32
  #exit
  gtpu-service SxC
    echo-interval 60
    echo-retransmission-timeout 1
    bind ipv4-address 50.50.17.4 ipv6-address 2001:4888:50:50::17 bearer-type all
  exit
  gtpu-service n9-ingress
    bind ipv4-address 60.60.17.5 bearer-type all
  exit
  gtpu-service pgw-gtpu
    bind ipv4-address 50.50.17.5 bearer-type all
  exit
  gtpu-service sgw-gtpu-egress
    bind ipv4-address 50.50.17.6 bearer-type all
  exit
  gtpu-service sgw-gtpu-ingress-phazr
    echo-interval 60
    echo-retransmission-timeout 1
    bind ipv4-address 60.60.17.6 bearer-type all
  exit
  sx-service sx-svc
    instance-type userplane
    bind ipv4-address 50.50.17.4 ipv6-address 2001:4888:50:50::17
    sxa max-retransmissions 5
    sxa retransmission-timeout-ms 1000
    sxb max-retransmissions 5
    sxb retransmission-timeout-ms 1000
    sxab max-retransmissions 5
    sxab retransmission-timeout-ms 1000
    sx-protocol heartbeat retransmission-timeout 15
    sx-protocol heartbeat max-retransmissions 3
    sx-protocol association debug-mode debug-reattempt-timeout 1
    no sx-protocol compression
  exit
  user-plane-service user_plane_svc
    associate gtpu-service pgw-gtpu upf-ingress
    associate gtpu-service n9-ingress upf-ingress interface-type n9
    associate gtpu-service sgw-gtpu-ingress-phazr sgw-ingress
    associate gtpu-service sgw-gtpu-egress sgw-egress
    associate gtpu-service SxC cp-tunnel
    associate sx-service sx-svc
    associate fast-path service
    associate control-plane-group SAEGW
    associate userplane-load-control-profile LCP
    associate userplane-overload-control-profile OLCP
  exit
```

The following is a sample configuration to separate the N3 and N9 BGP routes for vUPF.

```
context SAEGW
  router bgp 61441
    neighbor 20.14.41.100 remote-as 65000
    neighbor 20.14.41.100 ebgp-multihop
    neighbor 20.14.41.100 update-source 20.14.41.98
    neighbor 2001:4888:121:1441::1 remote-as 65000
    neighbor 2001:4888:121:1441::1 ebgp-multihop
    neighbor 2001:4888:121:1441::1 update-source 2001:4888:121:1441::2
```

```
  address-family ipv4
    redistribute connected
    redistribute static
  #exit
  address-family ipv6
    neighbor 2001:4888:121:1441::1 activate
    redistribute connected
    redistribute static
  #exit
#exit
interface SxC loopback
  ip address 50.50.18.4 255.255.255.255
#exit
interface SxC-ipv6 loopback
  ipv6 address 2001:4888:50:50::18/128
#exit
interface gn-1
  ip address 20.14.41.98 255.255.255.0
  ipv6 address 2001:4888:121:1441::2/64 secondary
#exit
interface gn-2
  ip address 20.14.42.98 255.255.255.0
  ipv6 address 2001:4888:121:1442::2/64 secondary
#exit
interface n9-egress loopback
  ip address 50.50.18.8 255.255.255.255
#exit
interface pgw-gtpu loopback
  ip address 50.50.18.5 255.255.255.255
#exit
interface pgw-gtpu-ipv6 loopback
  ipv6 address 2001:4888:50:50::18:5/128
#exit
interface sgw-gtpu-egress loopback
  ip address 50.50.18.6 255.255.255.255
#exit
interface sgw-gtpu-egress-ipv6 loopback
  ipv6 address 2001:4888:50:50::18:6/128
#exit
interface sgw-gtpu-ingress-phazr loopback
  ip address 50.50.18.7 255.255.255.255
#exit
interface sgw-gtpu-ingress-phazr-ipv6 loopback
  ipv6 address 2001:4888:50:50::18:7/128
#exit
subscriber default
exit
aaa group default
#exit
gtpp group default
  gtpp limit-secondary-rat-usage 32
#exit
gtpu-service SxC
  echo-interval 60
  echo-retransmission-timeout 1
  bind ipv4-address 50.50.18.4 ipv6-address 2001:4888:50:50::18 bearer-type all
exit
gtpu-service n9-egress
  bind ipv4-address 50.50.18.8 bearer-type all
exit
gtpu-service pgw-gtpu
  bind ipv4-address 50.50.18.5 ipv6-address 2001:4888:50:50::18:5 bearer-type all
exit
gtpu-service sgw-gtpu-egress
```

```
    bind ipv4-address 50.50.18.6 ipv6-address 2001:4888:50:50::18:6 bearer-type all
  exit
  gtpu-service sgw-gtpu-ingress-phazr
    echo-interval 60
    echo-retransmission-timeout 1
    bind ipv4-address 50.50.18.7 ipv6-address 2001:4888:50:50::18:7 bearer-type all
  exit
  sx-service sx-svc
    instance-type userplane
    bind ipv4-address 50.50.18.4 ipv6-address 2001:4888:50:50::18
    sxa max-retransmissions 5
    sxa retransmission-timeout-ms 1000
    sxb max-retransmissions 5
    sxb retransmission-timeout-ms 1000
    sxab max-retransmissions 5
    sxab retransmission-timeout-ms 1000
    sx-protocol heartbeat retransmission-timeout 15
    sx-protocol heartbeat max-retransmissions 3
    sx-protocol association debug-mode debug-reattempt-timeout 1
    no sx-protocol compression
  exit
  user-plane-service user_plane_svc
    associate gtpu-service pgw-gtpu upf-ingress
    associate gtpu-service n9-egress upf-egress
    associate gtpu-service sgw-gtpu-ingress-phazr sgw-ingress
    associate gtpu-service sgw-gtpu-egress sgw-egress
    associate gtpu-service SxC cp-tunnel
    associate sx-service sx-svc
    associate fast-path service
    associate control-plane-group SAEGW
    associate userplane-load-control-profile LCP
    associate userplane-overload-control-profile OLCP
  exit
  ip route 0.0.0.0 0.0.0.0 20.14.41.100 gn-1
  ip igmp profile default
  #exit
#exit
```

Note For RCM UPF, add all the IP prefix-lists in day 0.5 configuration.

Verifying the UPF Ingress Interface Type Configuration

To verify the configuration, use the following command:

- show user-plane-service all, on page 553

show user-plane-service all

The **show user-plane-service all** command displays user-plane service information.

The following is a sample output of this command:

```
show user-plane-service all
Service name                        : user-plane-service
  Service-Id                        : 7
  Context                           : EPC2-UP
  Status                            : NOT STARTED
  UPF Ingress GTPU Service          : sx-gtpu-service
```

```
UPF Ingress N3 Interface Type GTPU Service: n3-ingress
UPF Ingress N9 Interface Type GTPU Service: n9-ingress
UPF Ingress S5U Interface Type GTPU Service: s5u-ingress
UPF Ingress S8U Interface Type GTPU Service: s8u-ingress
UPF Egress GTPU Service           : Not defined
SGW Ingress GTPU Service          : SGW-Ingress
SGW Egress GTPU Service           : SGW-Egress
Control Plane Tunnel GTPU Service : control_gtpu
Sx Service                        : sxu
Control Plane Group               : g1
Fast-Path service                 : Disabled
```

NOTES:

- The UPF configures only one of the interface types, **pgw-ingress** or **upf-ingress** in a single user plane service.

Multiple GTPU IP Addresses Support

Table 118: Feature History

Feature Name	Release Information	Feature Description
Binding Multiple GTPU IP Addresses	2023.04	UPF allows binding of multiple GTPU IP addresses to provide high uplink throughput in Private 5G deployments. UPF creates unique GTPU 3-tuple hash entries to ensure uniform distribution of ingress traffic on all VPP worker threads.

Feature Description

In the Private 5G networks, where the uplink traffic is more than the downlink traffic, the number of gNodeBs are limited. UPF achieves the optimal distribution of subscriber traffic using the Receive Side Scaling (RSS) concept. However, due to limited number of VPP worker threads on UPF, the ingress traffic does not get evenly distributed across them. It limits the UPF from optimally utilizing all the worker threads, and therefore impacts the subscriber's uplink throughput.

UPF allows utilizing the VPP worker threads in an optimal manner by creating the unique GTPU 3-tuple hash entries on the UPF's forwarding plane. It ensures an optimum distribution of ingress traffic on the N3 interface. It enables UPF to perform traffic processing to its maximum capacity ensuring high uplink throughput for Private 5G networks.

How it Works

The even distribution of ingress traffic across all the available worker threads on the UPF is ensured by creating the unique 3-tuple hash entries.

The GTPU header source IP is the gNodeB's N3 service egress IP, while the source port is the standard port in the absence of port randomization support on the gNodeB. With these two values being constants per gNodeB, the only variable that can be achieved in the GTPU header is the destination IP (N3 service ingress IP address).

The following points describe how the unique 3-tuple hash entries are created:

- GTPU addresses are assigned in FTEID in Sx Session Establish Response for a subscriber session in a round-robin manner. However, this round-robin mechanism is ensured only within the session manager.
- When a session manager is selected, the GTPU address in FTEID is allocated in a round-robin manner to all the sessions landing on the selected session manager.
- Deletion of IP address from the GTPU service or the interface does not cause the GTPU service to stop. The existing calls will be deleted. However, the uplink traffic of the subscriber to which the deleted IP address was assigned will not be received at UPF.
- If every IP address is configured under a separate interface, the interface must be configured as a loopback (multiple non-loopback interfaces cannot share the same subnet). If any interface goes down and comes alive again, the GTPU service does not stop as the interfaces are configured as loopback.
- Every address pair is one GTPU endpoint for a GTPU service. Every endpoint should belong to the same subnet.



Note

- For optimum performance, the number of GTPU bind addresses must be equal to the number of VPP worker threads.
- A maximum of 12 IP address pairs (either IPv4 or IPv6 or IPv4+IPv6) can be configured per GTPU service.

Limitations

This feature has the following known limitations::

- If the GTPU Echo is enabled and the first bind IP address is deleted, the Echo request from the remote peer will not reach the UPF causing the remote peer to detect path failure. The subsequent downlink traffic from the UPF for that subscriber results in an error indication received at UPF.

  The GTPU services stop only if all the bind IP addresses are deleted. These IP addresses can be under single interface or separate interfaces.
- UPF does not allow to configure different IP address types to be configured per GTPU service i.e., IPv4 + IPv6 or IPv4 + IPv4/IPv6 combinations are not allowed. Only the combinations of IPv4 or IPv6 or IPv4+IPv6, are allowed to be configured.

Configuring Multiple Bind Addresses on N3 Interface

To configure multiple bind addresses on the N3 interface, use the following sample configuration:

```
config
  context context_name
    gtpu-service service_name
      bind { ipv4-address ipv4_address | ipv6-address ipv6_address [
bearer-type { non-ims-media | ims-media | all } ] }
      exit
```

NOTES:

- **bind { ipv4-address** *ipv4_address* **| ipv6-address** *ipv6_address* **}**—Bind the GTPU service to the IPv4 address or IPv6 address on the N3 interface.
- **[bearer-type { non-ims-media | ims-media | all }]**—Specify the type of bearer to be associated with the bind address. **bearer-type** allows three types of media traffic:
  - **non-ims-media**—Configure bind address for non-ims media only.
  - **ims-media**—Configure bind address for ims-media traffic only.
  - **all**—Configure bind address to handle all types of bearer traffic. This is the default behavior.



Note If GTPU echo is enabled in GTPU service, the Echo request will be started with the first bind address configured in GTPU service. The bind address configuration is stored in the sorted list. As a result, the GTPU echo starts with the numerically lowest bind IP address.

Configuration Example

The following is an example configuration for binding multiple GTPU IP addresses on the N3 interface:

```
config
context EPC2-UP
cups enable
interface uplane_interface
    ip address <ip_address 1> <subnet mask>
    ip address <ip_address 2> <subnet mask> secondary
    ip address <ip_address 3> <subnet mask> secondary
    ip address <ip_address 4> <subnet mask> secondary
    ip address <ip_address 5> <subnet mask> secondary
    ip address <ip_address 6> <subnet mask> secondary
    ip address <ip_address 7> <subnet mask> secondary
#exit
gtpu-service N3U-service
    bind ipv4-address <ip_address 1> bearer-type all
    bind ipv4-address <ip_address 2> bearer-type all
    bind ipv4-address <ip_address 3> bearer-type all
    bind ipv4-address <ip_address 4> bearer-type all
    bind ipv4-address <ip_address 5> bearer-type all
    bind ipv4-address <ip_address 6> bearer-type all
exit
```

Configuration Verification

To verify the configuration, use the following command:

show gtpu-service name N3U-service

The following is a sample output of this command:

```
show gtpu-service name N3U-service

Service name:             N3U-service
  Context:                  EPC2-UP
  State:                    Started
  Echo Interval:            Disabled
```

```
Sequence number:          Disabled
Include UDP Port Ext Hdr: FALSE
Max-retransmissions:         4
Retransmission Timeout:   5 (secs)
IPSEC Tunnel Idle Timeout:60 (secs)
GTPU IP QOS DSCP value:10
Allow Error-Indication:   Disabled
Address List:
               <ip_address 1>      all
               <ip_address 2>      all
               <ip_address 3>      all
               <ip_address 4>      all
               <ip_address 5>      all
               <ip_address 6>      all
               <ip_address 7>      all
GTPU UDP Checksum:  Enabled - Attempt Optimize Default Mode
Path Failure Detection
on gtp echo msgs:         Set
Path Failure Clear Trap : non-echo
GTPU Source-port Configuration : Use Standard Port
  GTPU Source-port Offset      : N/A
GTPU Peer Statistics Threshold : 16000
    1
```

OAM Support

This section describes operations, administration, and maintenance support for this feature.

Monitoring Support

This section describes feature-level monitoring support information.

show user-plane-service gtpu local-addresses

This command displays the number of active bearers/sessions using specific bind address.

For Single Converged 4G Call:

```
show user-plane-service gtpu local-addresses
  Service name:         N3U-service
  Address List:                                                    GTPv1 Sessions:
     <ip_address 1>                                                    1
  Service name:         SGW-Egress
  Address List:                                                    GTPv1 Sessions:
     <ip_address 2>                                                    1
  Service name:         SGW-Ingress
Address List:                                                   GTPv1 Sessions:
     <ip_address 3>                                                    1
```

Note

- In 4G calls, there will be one GTPU session per bearer.
- This CLI output will not display the total number of unique subscribers per GTPU address.

For 5G Call:

```
show user-plane-service gtpu local-addresses
  Service name:         N3U-service
  Address List:                                                    GTPv1 Sessions:
```

```
    <ip_address 1>                                                    1
  Service name:         SGW-Egress
  Address List:                                                GTPv1 Sessions:
No active sessions exists
  Service name:         SGW-Ingress
  Address List:                                                GTPv1 Sessions:
No active sessions exists
```

show user-plane-service gtpu statistics local-address

This command displays the statistics per GTPU bind address.

```
show user-plane-service gtpu statistics local-address < Local IPv4/IPv6 Address>
Session Stats:
    Current:                                  1
    Current (IMS-media):                      0
    Total Setup:                              2
    Total Setup (IMS-media):                  0
    Current gtpu v0 sessions:                 0
    Current gtpu v1 sessions:                 2

  Total Data Stats:
    Uplink Packets:                   0  Uplink Bytes:               0
    Downlink Packets:                 0  Downlink Bytes:             0
    Packets Discarded:                0  Bytes Discarded:            0
    Uplink Packets (IMS-media):                0  Uplink Bytes (IMS-media):
    0
    Downlink Packets (IMS-media):              0  Downlink Bytes (IMS-media):
    0
    Packets Discarded (IMS-media):             0  Bytes Discarded (IMS-media):
    0

 QoS Stats:

    QCI 1:
         Uplink Packets:              0  Uplink Bytes:               0
         Downlink Packets:            0  Downlink Bytes:             0
         Packets Discarded:           0  Bytes Discarded:            0

    QCI 2:
         Uplink Packets:              0  Uplink Bytes:               0
         Downlink Packets:            0  Downlink Bytes:             0
         Packets Discarded:           0  Bytes Discarded:            0

    QCI 3:
         Uplink Packets:              0  Uplink Bytes:               0
         Downlink Packets:            0  Downlink Bytes:             0
         Packets Discarded:           0  Bytes Discarded:            0

    QCI 4:
         Uplink Packets:              0  Uplink Bytes:               0
         Downlink Packets:            0  Downlink Bytes:             0
         Packets Discarded:           0  Bytes Discarded:            0

    QCI 5:
         Uplink Packets:              0  Uplink Bytes:               0
         Downlink Packets:            0  Downlink Bytes:             0
         Packets Discarded:           0  Bytes Discarded:            0

    QCI 6:
         Uplink Packets:              0  Uplink Bytes:               0
         Downlink Packets:            0  Downlink Bytes:             0
         Packets Discarded:           0  Bytes Discarded:            0
```

```
QCI 7:
      Uplink Packets:             0  Uplink Bytes:               0
      Downlink Packets:           0  Downlink Bytes:             0
      Packets Discarded:          0  Bytes Discarded:            0

 QCI 8:
      Uplink Packets:             0  Uplink Bytes:               0
      Downlink Packets:           0  Downlink Bytes:             0
      Packets Discarded:          0  Bytes Discarded:            0

 QCI 9:
      Uplink Packets:             0  Uplink Bytes:               0
      Downlink Packets:           0  Downlink Bytes:             0
      Packets Discarded:          0  Bytes Discarded:            0

 QCI 65:
      Uplink Packets:             0  Uplink Bytes:               0
      Downlink Packets:           0  Downlink Bytes:             0
      Packets Discarded:          0  Bytes Discarded:            0

 QCI 66:
      Uplink Packets:             0  Uplink Bytes:               0
      Downlink Packets:           0  Downlink Bytes:             0
      Packets Discarded:          0  Bytes Discarded:            0

 QCI 69:
      Uplink Packets:             0  Uplink Bytes:               0
      Downlink Packets:           0  Downlink Bytes:             0
      Packets Discarded:          0  Bytes Discarded:            0

  QCI 70:
      Uplink Packets:             0  Uplink Bytes:               0
      Downlink Packets:           0  Downlink Bytes:             0
      Packets Discarded:          0  Bytes Discarded:            0

 QCI 80:
      Uplink Packets:             0  Uplink Bytes:               0
      Downlink Packets:           0  Downlink Bytes:             0
      Packets Discarded:          0  Bytes Discarded:            0

 QCI 82:
      Uplink Packets:             0  Uplink Bytes:               0
      Downlink Packets:           0  Downlink Bytes:             0
      Packets Discarded:          0  Bytes Discarded:            0

 QCI 83:
      Uplink Packets:             0  Uplink Bytes:               0
      Downlink Packets:           0  Downlink Bytes:             0
      Packets Discarded:          0  Bytes Discarded:            0

   Non-Std QCI(Non-GBR):
       Uplink Packets:             0  Uplink Bytes:                 0
       Downlink Packets:           0  Downlink Bytes:         0
       Packets Discarded:          0  Bytes Discarded:          0
       Uplink Packets:             0  Uplink Bytes:                 0
       Downlink Packets:           0  Downlink Bytes:         0
       Packets Discarded:          0  Bytes Discarded:          0

   Total uplink packets GBR QCI's:               0
   Total uplink Bytes GBR QCI's:                 0
   Total Downlink packets GBR QCI's:             0
   Total Downlink Bytes GBR QCI's:               0
   Total uplink packets Non-GBR QCI's:           0
   Total uplink Bytes Non-GBR QCI's:             0
```

```
Total Downlink packets Non-GBR QCI's:          0
Total Downlink Bytes Non-GBR QCI's:            0

 Path Management Messages:
   Echo Request Rx:                  0  Echo Response Rx:            0
   Echo Request Tx:                  0  Echo Response Tx:            0
   SuppExtnHdr  Tx:                  0  SuppExtnHdr   Rx:            0

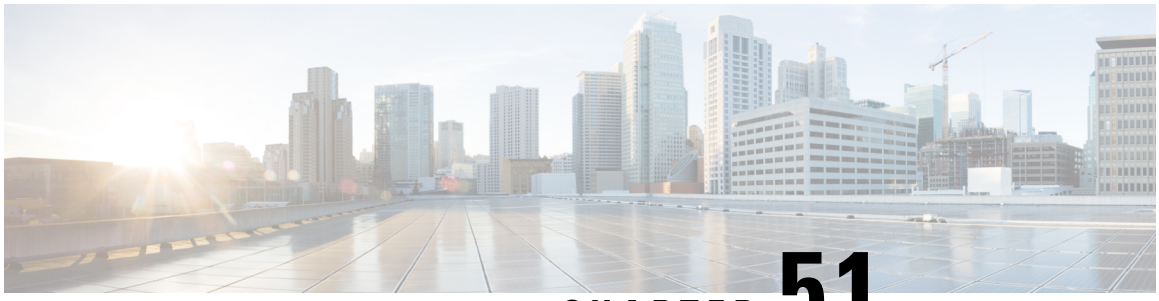
 Peer Stats:
   Total GTPU Peers:                             1
   Total GTPU Peers with Stats:                  1

Tunnel Management Messages:
    Error Indication Tx:                          0
    Error Indication Rx:                          0
    Error Indication Rx Discarded:                0
```

show npu utilization table

This command displays the VPP worker thread CPU utilization:

```
[local]UPF1# show npu utilization table
           ---------vpp--------
   lcore     now   5min  15min
--------  ------ ------ ------
 01/0/01     15%    11%    12%
 01/0/02     15%    13%    14%
 01/0/03     23%    23%    24%
 01/0/04     19%    21%    20%
 01/0/05     22%    20%    20%
 01/0/06     12%    11%    11%
 01/0/07     16%    16%    16%
 01/0/08     15%    17%    17%
 01/0/09     11%    12%    12%
```



CHAPTER 51

UPF Local Configuration

This chapter covers the following topics:

- Feature Summary and Revision History, on page 561
- Feature Description, on page 562
- Configuring the Local Configuration Support for UPF, on page 563
- TACACS+ TLS support, on page 564

Feature Summary and Revision History

Summary Data

Table 119: Summary Data

Applicable Product(s) or Functional Area	5G-UPF
Applicable Platform(s)	VPC-SI SMI
Feature Default Setting	Disabled – Configuration Required
Related Changes in this Release	Not Applicable
Related Documentation	Not Applicable

Revision History

Table 120: Revision History

Revision Details	Release
First Introduced.	2020.02.0

Feature Description

The support for processing static and predefined rules in Control and User Plane Separation of EPC nodes (CUPS) architecture is dependent on the ruledef, rulebase, and charging action. For processing L3/L4 static and predefined rules, this information is made available at the control-plane in CUPS architecture. The control plane sends all these information to the associated user-plane using the PFD management message. The UPF cannot use the PFD management message to work with CN-SNF. With this feature, the local configuration support for the User Plane Function (UPF) is enabled, which allows the UPF to work with CN-SNF.

How it Works

The Access Control System (ACS) command line interface (CLI) is configured on the user-plane and the CLI module sends it to the ACS Controller (ACSCtrl). The ACSCtrl verifies the CLI and sends it to the Session Controller (SessCtrl). The SessCtrl stores the configuration in the SCT.

The SessCtrl maintains and stores different configuration types in a skiplist. When the length of the skiplist reaches the maximum (BULK configuration length) for a particular configuration type, the entire list is pushed in BULK from the Sessctrl to the Session Manager (SessMgr). As a result, the number of messenger event/message transactions between proclets is greatly reduced since the configurations are sent in BULK in a single message. On the expiry of the bulk configuration timeout (2 seconds), the Bulk Configuration timer – which runs constantly at the Session Controller – pushes the different types of configurations to the SessMgrs.

- The following configuration types are supported for the Bulk Configuration push:
- Ruledef
- Charging Action
- Action Priority Lines
- Group of Ruledef Configuration
- Rule in Group of Ruledef Configuration
- Rulebase L3/L4/L7 Info Configuration
- APN Configuration
- ACS service Configuration

The configurations are pushed only through the bulk push mechanism for configurations that are either added or modified. On the other hand, when configurations are deleted, it is removed immediately without waiting for any response from the Bulk configuration push timer. The deleted configuration is removed from the SCT and other SessMgrs immediately.



Note The Bulk configuration timeout function is invoked forcefully to push all the pending configurations to the SessMgrs before pushing the configuration delete to avoid any race conditions.

- The configuration changes applied to all the new and existing calls are listed in Table as follows.

Table 121: Configuration Changes on New and Existing Call Flows

Change in Configuration	Impact on Existing Calls Current Flows	Impact on Existing Calls New Flows	Impact on New Calls
Existing ruledef contents/New rule addition	Rule match is not enforced on existing flows after configuration change. TRM is not disengaged on existing flows. This may lead to billing issues if ruledef contents were changed for ongoing flows.	The configuration changes apply on new flows. For new flows, anyways fresh rule match would happen and the ruledef changes are applied on new flows for existing calls.	The configuration changes apply on new calls. For new flows, anyways fresh rule match would happen and the ruledef changes are applied on flows for new calls.
No Ruledef	Rule in use cannot be deleted.	Rule in use cannot be deleted	Rule in use cannot be deleted
New Group of Ruledefs/Changes to existing Group of Ruledefs contents (Add or Delete Rule in Group of Ruledefs)	Rule match is not enforced on existing flows after configuration change. TRM is not disengaged on existing flows. This may lead to billing issues if Group of Ruledefs contents were changed for ongoing flows.	The configuration changes apply on new flows. For new flows, anyways fresh rule match would happen and the Group of Ruledefs changes are applied on new flows for existing calls.	The configuration changes apply on new calls. For new flows, anyways fresh rule match happens and the Group of Ruledefs changes are applied on flows for new calls.
No Group of Ruledefs	Group of Ruledefs in use cannot be deleted	Group of Ruledefs in use cannot be deleted	Group of Ruledefs in use cannot be deleted
No Rule in GoR	Flows continue to match the ruledef defined in Group of Ruledefs unless the ruledef itself is deleted	New flows go through a fresh rule match and configuration change takes effect.	New flows go through a fresh rule match and configuration change takes effect.
Action Priority Changes/Action Priority addition	TRM is not disengaged for ongoing flows. configuration changes do not apply on existing flows	Configuration changes apply on new flows.	Configuration changes apply on new calls.
No Action Priority	No Impact on existing flows	Configuration changes apply on new flows.	Configuration changes apply on new calls.
Rulebase parameters change	Some parameter changes apply on existing calls	Some parameter changes apply on existing calls	Configuration changes apply on new calls
No Rulebase	No Rulebase is not supported	No Rulebase is not supported	No Rulebase is not supported
No APN	No APN is not supported	No APN is not supported	No APN is not supported
IP source violation	No impact on existing calls	No impact on existing calls	Configuration changes apply on new calls.

Configuring the Local Configuration Support for UPF

Use the following CLI commands to configure the User Plane Function (UPF) locally.

```
configure
  require upf
  end
```

TACACS+ TLS support

Table 122: Feature history

Feature name	Release	Feature dscription
Added TACACS+ TLS support	2026.03.0	StarOS now supports TLS-secured TACACS+ communication, enabling encrypted AAA traffic and certificate-based trust validation between StarOS nodes and TACACS+ servers. This enhancement strengthens management-plane security while maintaining compatibility with existing TACACS+ deployments. TLS support must be enabled on both the StarOS node and the TACACS+ server to establish secure communication.

TACACS+ TLS support is an optional security enhancement, disabled by default, that integrates TLS 1.3 for strong mutual authentication and encryption. Administrators may enable this feature to secure administrative traffic, effectively replacing insecure MD5 obfuscation methods with robust TLS protection

TACACS+ TLS support is a security enhancement that:

- Establishes a TLS 1.3 session to the TACACS+ server,
- Enables certificate-based mutual authentication (mTLS) between the client and server, and
- Replaces insecure MD5 obfuscation methods with robust TLS protection.

Policy-driven TLS management allows network administrators to manage certificates via the platform secret store without requiring service restarts.

Configure TACACS+ TLS

Enable TLS transport for TACACS+ to secure administrative traffic.

Before you begin

- Ensure the `.pem` files for the client and CA certificates are accessible via the configured URL.
- Verify that the TACACS+ server supports TLS 1.3.

Follow these steps to configure TACACS+ TLS:

Procedure

Step 1 Configure the client certificate and private key for mutual TLS (mTLS) authentication with the TACACS+ server

```
[local]# config
[local](config)# certificate name tls-client pem url /sftp/client_v3.crt private-key pem url
/sftp/client_v3.key
[local](config)# end
```

Step 2 Configure the CA certificate to verify the TACACS+ server certificate during the TLS handshake

```
[local]# config
[local](config)# certificate name tls-ca pem url /sftp/ca.pem private-key pem url /sftp/ca.key
[local](config)# end
```

Step 3 Associate the client certificate profile with the TACACS+ server for TLS connections. The profile name must match the certificate name configured in Step 1.

```
[local](config)# tacacs mode
[local](config-tacacs)# server priority 1 ip-address 10.88.0.1 port <6049> password <secret>
[local](config-tacacs)# tls-profile <profile-name>
```

Step 4 Associate the CA certificate profile with the TACACS+ server for server certificate verification. The profile name must match the CA certificate name configured in Step 2.

```
[local](config-tacacs)# tls-ca-profile <ca-profile-name>
```

Step 5 (Optional) Enable hostname/IP verification to prevent man-in-the-middle attacks by verifying that the server certificate Subject Alternative Name (SAN) matches the configured server IP.

```
[local](config-tacacs)# tls-verify-hostname
```

To disable hostname verification (default):

```
[local](config-tacacs)# no tls-verify-hostname
```

Step 6 Complete the TACACS+ configuration.

```
[local](config-tacacs)# exit
[local](config)# aaa tacacs+
[local](config)# end
```

The TACACS+ server is configured to use the TLS port (6049) instead of the cleartext port (49), and the TACACS+ connection is established over a secure TLS 1.3 tunnel.

TACACS+ TLS monitoring and troubleshooting

This section provides the commands and field descriptions required to monitor TLS-enabled TACACS+ sessions.

Use the `show tacacs session statistics` command to verify the status of your TLS-enabled sessions. The output confirms if TLS is active for the current session, displaying the associated profiles or 'N/A' if the feature is disabled.

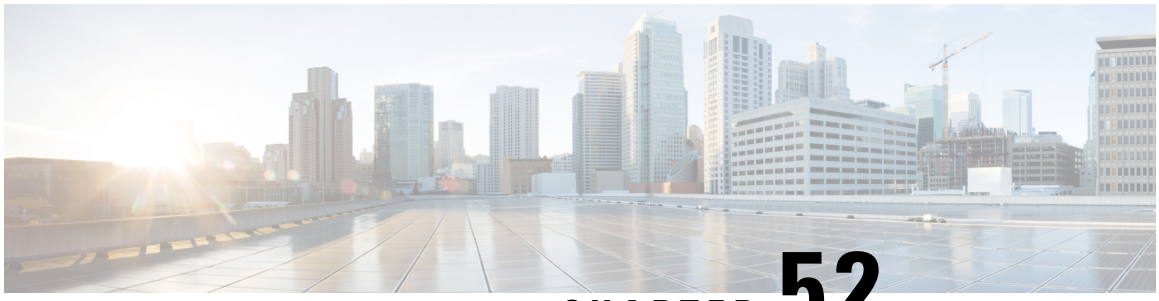
Output Examples

When TACACS+/TLS is enabled:

```
tls enabled                    : yes
tls profile                    : tls-client
tls ca profile                 : tls-ca
```

When TACACS+/TLS is disabled:

```
tls enabled                    : no
tls profile                    : N/A
tls ca profile                 : N/A
```



CHAPTER 52

Load and Overload Control Over N4/Sx Interface

- Feature Summary and Revision History, on page 567
- Feature Description, on page 568
- How Load and Overload Control Works, on page 569
- Enable Load and Overload Control on N4 or Sx Interface, on page 572
- Monitoring and Troubleshooting, on page 575

Feature Summary and Revision History

Summary Data

Applicable Product(s) or Functional Area	UPF
Applicable Platform(s)	VPC-SI SMI
Feature Default Setting	Disabled – Configuration Required
Related Changes in this Release	Not Applicable
Related Documentation	Not Applicable

Revision History

Table 123: Feature History

Feature Name	Release Information	Description
Load and Overload Control over N4/Sx Interface	2024.04.1	This feature allows UPF to enable load and overload control mechanisms to handle the messages during overloaded or self-protection state. Commands Introduced: • **upf-load-control-profile** *profile_name* • **upf-overload-control-profile** *profile_name* • **associate { upf-load-control-profile** *profile_name* **| upf-overload-control-profile** *profile_name* **}** **Default Setting:** Disabled-Configuration required to enable.
Reporting of Load Control over N4 Interface	2020.02.0	This feature enables load control on UPF. This feature allows UPF to send its load information to the SMF in order to balance PFCP session load across the UPFs as per their effective load. **Default Setting:** Disabled-Configuration required to enable.

Feature Description

There are three states of UPF and their default values at any given time:

- **Loaded State**: When the load on UPF is up to 70%.
- **Overloaded State**: When the load on UPF is between 70–85%.
- **Self-Protection State**: When the load on UPF is between 85–100%.

Note These default values can be configured.

Load control enables the UPF to send its load information to the CP-F in order to balance PFCP session load across the UPF according to their effective load. The load information reflects the operating status of the resources of the UPF. Load control allows for better balancing of the PFCP session load to prevent overload.

Overload control enables throttling of new session requests toward a particular UPF. In the overloaded and self-protection state, if the load exceeds the configured thresholds, the session deletion process is initiated for non-priority sessions.

NOTE: Overload mitigation actions are not triggered even if the UPF reports a high load.

How Load and Overload Control Works

CP-F sends the load and overload features as supported in the Sx Association Request and Response message to UPF. Based on this information, UPF decides whether to send the load or overload information toward CP-F peer.

Load or overload support at CP-F is configured as part of the Sx Service node configuration. This information is sent to the UP during Sx Association Response or Sx Association Update request, if the information has changed through dynamic configuration.

When the UPF is in the overload state, the OCI IE is sent to the control-plane function with reduction metric.

Note

- If the load or overload control feature is enabled or disabled on the CP-F node, the CP-F node shall initiate an Association Update Request toward UPF to inform the updated support.
- UPF does not take any action when the cause code **PFCP_ENTITY_IN_CONGESTION** cause is received in any response message.

Load and Overload Factor Calculation

Load Factor Calculation

The load factor represents the total load on the system at any given time. The UPF is considered as loaded if either of its CPU, Memory, or Session count is high. Therefore, the load factor is calculated based on system parameters like System CPU, System Memory, Session Count.

Overload Factor Calculation

The UPF is considered as overloaded if the load exceeds any of its configured threshold values. The following parameters are considered for calculating the overload factor in the self-protection mode:

System Level Parameters: CPU, Memory, Session Count

VPP: VPP CPU

Load and Overload Reporting to CP-F

UPF sends load or overload information to CP-F through N4 messages.

If the UPF is in...	It sends...
Load State	LCI IE
Overload State	OCI IE

CP-F parses all the incoming N4 messages and N4 responses for LCI and OCI IEs. Whenever CP-F receives LCI IE or OCI IE from a peer node, CP-F learns about the load and overload parameter of that UPF node.

For more details on how CP-F node handles the LCI and OCI IEs respectively, see the following guides:

- *UCC 5G Session Management Function Configuration and Administration Guide, Release 2024.04*.
- *UCC 5G cnSGWc Configuration and Administration Guide, Release 2024.04*.

Load Information Reporting to CP-F

The UPF sends its load control information to reflect the operating status of its resources at the node level. It allows the CP-F to use this information to augment the UPF selection procedures. The load control information is piggybacked in PFCP request or response messages such that the exchange of load control information does not trigger extra signaling.

Load Control Information IE can be present in the following messages from UPF to CP-F:

- PFCP Session Establishment Response
- PFCP Session Modification Response
- PFCP Session Deletion Response
- PFCP Session Report Request

Information Elements of LCI IE

Load information is reported using the following elements present in the LCI IE:

- **Load Control Sequence Number**: This number is incremented whenever the updated or new load control information is sent. The parameter **current-time** is used for Load Control Sequence Number.
- **Load Metric**: This parameter indicates the current load level of the UPF. It is sent as a percentage within the range of 0-100.



Note

- Multiple SessMgrs report the same value of Load Metric with the same sequence number.
- The System timestamp is used as a Sequence Number.

Overload Information Reporting to CP-F

When the UPF is in an overload state, it sends OCI IE in the N4/Sx messages. The OCI IE can be present in the following messages:

- PFCP Session Establishment Response
- PFCP Session Modification Response
- PFCP Session Deletion Response
- PFCP Session Report Request

Information Elements of OCI IE

Overload information is reported using the following elements present in the OCI IE:

- **Overload Control Sequence Number**: This number is incremented whenever the updated or new overload control information is sent. The parameter **current-time** is used for Overload Control Sequence Number.
- **Overload Reduction Metric**: In the self-protection mode, the CP function throttles the Sx Establishment request message toward UPF. The messages are throttled based on the reported Overload Reduction metric value. It is sent as a percentage within the range of 0–100.
- **Period of Validity**: This parameter indicates the length of time during which the overload condition specified by the OCI IE is to be considered as valid (unless overridden by subsequent new overload control information).
- **Overload Control Information Flags**: This parameter contains the flags for overload control-related information. The Bit 1 represents Associate OCI with Node ID (AOCI). This flag is set to "1" by the UP function, if the OCI IE is included in the PFCP Session Establishment Response and the Cause IE is set to a rejection cause value.

Session Clear in Self-Protection Mode

If the system is in the self-protection mode and the system state keeps deteriorating, it may lead to session throttling at UPF.

The session reports are sent in following two ways:

- If the EPFAR feature is supported or negotiated on the peers, the UPF initiates the PFCP Session Report Request message for the sessions with report type as 'USAR' or 'UISR'. (When there is a non-zero usage report for the PFCP) or UISR (UP Initiated Session Request) if there is no usage report to send.
  - **USAR (Usage Report)**: When there is a non-zero usage report for the PFCP.
  - **USIR (UP Initiated Session Request)**: When there is no usage report to send.

  The 'cause' IE in PFCP Session Report Request will be set to **PFCP_CAUSE_SELF_PROTECTION_TERMINATION (205)**. The session is deleted after sending a Session Report Request.
- If the EPFAR feature is not negotiated with the peers, it initiates the PFCP Session Report Request message for the sessions with report type as 'SPTER' and waits for the peer to initiate the session deletion gracefully. If the session deletion is not initiated within 120 sec, the session will be deleted locally.

Enable Load and Overload Control on N4 or Sx Interface

Follow these steps to enable load and overload control on UPF:

- Configure Load Control Profiles
- Configure Overload Control Profiles
- Associate the Load and Overload Control Profiles with User Plane Service

Configure the Load Control Profile

Follow these steps to configure the load control profile:

Procedure

Step 1 Configure a load control profile instance in the Global Configuration mode using the command **upf-load-control-profile** *profile_name*.

Example:

```
local]UPlane# config
[local]UPlane(config)# upf-load-control-profile loadcontrolprof
[local]UPlane(config-upf-load-control-profile)#
```

Step 2 Configure advertisement interval and change factor for under the configured load control profile using the commands **inclusion-frequency advertisement-interval** *interval_value* and **inclusion-frequency change-factor** *changefactor_value*.

Example:

```
[local]UPlane(config-upf-load-control-profile)# inclusion-frequency advertisement-interval
 300
 [local]UPlane(config-upf-load-control-profile)# inclusion-frequency change-factor 5
```

The parameter **inclusion-frequency** specifies the frequency at which the LCI IE is sent out. The following 2 factors are considered for sending LCI IE:

- The parameter **advertisement-interval** is the interval at which LCI IE is sent out. The default value is 300 seconds.
- The parameter **change-factor** is the % change from the previous LCI factor, at which the LCI IE is sent out. The default value is 5%.

Step 3 Exit the current configuration mode using the command **end**.

Example:

```
[local]UPlane(config-upf-load-control-profile)# end
[local]UPlane# config#
```

Configure the Overload Control Profile

Follow these steps to configure the overload control profile:

Procedure

Step 1 Configure a load control profile instance in the Global Configuration mode using the command **upf-overload-control-profile** *profile_name*.

Example:

```
[local]UPlane# config
[local]UPlane(config)# upf-overload-control-profile ovrloadcontrolprof
[local]UPlane(config-upf-overload-control-profile)#
```

Step 2 Configure lower and upper limits of the overload threshold parameters (System level and VPP CPU) under the overload control profile.

Example:

```
[local]UPlane(config-upf-overload-control-profile)# overload-threshold system cpu lower
limit 80 upper-limit 90
[local]UPlane(config-upf-overload-control-profile)# overload-threshold system memory lower
 limit 70 upper-limit 80
[local]UPlane(config-upf-overload-control-profile)# overload-threshold system session-count
 lower limit 80 upper-limit 90
[local]UPlane(config-upf-overload-control-profile)# overload-threshold vpp-cpu lower limit
 50 upper-limit 80
```

The parameter **system** configures the overload system thresholds. System moves to the overload state once the lower-limit is crossed and moves to self-protection state after the upper-limit is crossed. The system state can change due to either system level or VPP CPU parameters.

Overload Control Parameters	Lower Limit (Overload Control State)	Upper Limit (Self-Protection State)
System CPU	80	90
System Memory	70	80
System Session Count	80	90
VPP CPU	50	80

Note
When the system enters the self-protection state, the Reduction Metric is set to 100%.

Step 3 Configure advertisement interval, change factor, and validity period under the configured overload control profile using the commands **inclusion-frequency advertisement-interval** *interval_value*, **inclusion-frequency change-factor** *changefactor_value*, and **validity-period** *validity_period*.

Example:

```
[local]UPlane(config-upf-overload-control-profile)# inclusion-frequency
advertisement-interval 300
```

```
[local]UPlane(config-upf-overload-control-profile)# inclusion-frequency change-factor 5
[local]UPlane(config-upf-overload-control-profile)# validity-period 300
```

The parameter **inclusion-frequency** specifies the frequency at which the OCI IE is sent out. The following 2 factors are considered for sending OCI IE:

- The parameter **advertisement-interval** is the interval at which OCI IE is sent out. The default value is 300 seconds.
- The parameter **change-factor** is the % change from the previous OCI factor, at which the OCI IE is sent out. The default value is 5%.
- The parameter **validity-period** is the validity of the OCI factor. The default value is 300.

Step 4 Exit the current configuration mode using the command **end**.

Example:

```
[local]UPlane(config-upf-overload-control-profile)# end
[local]UPlane(config)#
```

Associate the Load and Overload Control Profiles with User Plane Service

Follow these steps to associate the load and overload control profiles with the user plane service:

Before you begin

Before associating the load and overload control profile with the user plane service, you need to configure the following:

- Load Control Profile
- Overload Control Profile

Procedure

Step 1 Configure a context instance in the Global Configuration mode using the command **context** *context_name*.

Example:

```
[local]UPlane# config
[local]UPlane (config)# context EPC
[EPC]UPlane (config-ctx)#
```

Step 2 Configure a User Plane Service instance under the configured context using the command **user-plane-service** *userplane_service_name*.

Example:

```
[EPC]UPlane (config-ctx)# user-plane-service UPlane1
[EPC]UPlane (config-user-plane-service)#
```

Step 3 Configure the maximum number of sessions that are to be supported on the UPF using the command **load-control capacity** *session_value*.

Example:

```
[EPC]UPlane (config-user-plane-service)# load-control capacity 120000
[EPC]UPlane (config-user-plane-service)#
```

Step 4 Associate the load and overload control profiles with the configured user plane service using the command **associate { upf-load-control-profile** *profile_name* **| upf-overload-control-profile** *profile_name* **}**.

Example:

```
[EPC]UPlane (config-user-plane-service)# associate upf-load-control-profile loadcontrolprof
[EPC]UPlane (config-user-plane-service)# associate upf-overload-control-profile
ovrloadcontrolprof
[EPC]UPlane (config-user-plane-service)#
```

Step 5 Exit the current configuration mode using the command **end**.

Example:

```
[EPC]UPlane (config-user-plane-service)# end
[EPC]UPlane (config-ctx)#
```

Monitoring and Troubleshooting

This section provides information regarding the CLI commands available for monitoring and troubleshooting the feature.

Show Commands and Outputs

This section provides information regarding show commands and their outputs in support of this feature.

show upf-load-control-profile

The show command **show upf-load-control-profile [name** *name* **| all]** displays the load control profile for UPF.

```
[local]qvpc-si# show upf-load-control-profile name lci1
UPF Load Control Profile
--------------------------------
lci1
```

```
[local]qvpc-si# show upf-load-control-profile full name lci1
UPF Load Control Profiles
--------------------------------
UPF Load Control Profile Name: lci1
 Inclusion Frequency:
   Change Factor                        : 6
   Advertisement Interval               : 300
```

show upf-overload-control-profile

The show command **show upf-overload-control-profile [name** *name* **| all]** displays the overload control profile for UPF.

```
[local]qvpc-si# show upf-overload-control-profile name oci1
UPF Overload Control Profile
```

```
--------------------------------
oci1

[local]qvpc-si# show upf-overload-control-profile full name oci1
UPF Overload Control Profiles
--------------------------------

UPF Overload Control Profile Name: OC1
 System Thresholds:
   CPU Lower Limit           : 80
   CPU Upper Limit           : 90
   Memory Lower Limit        : 70
   Memory Upper Limit        : 80
   Session Count Lower Limit : 80
   Session Count Upper Limit : 90
 VPP Thresholds:
   CPU Lower Limit           : 50
   CPU Upper Limit           : 80
 Inclusion Frequency:
   Change Factor             : 5
   Advertisement Interval    : 300
 Validity Period             : 300
```

show user-plane-service statistics name

The show command **show user-plane-service statistics name** *service_name* displays the load and overload control factors along with other existing stats.

```
[local]asr5000# show user-plane-service statistics name <user_plane_service_name>
  UPF Overload Stats:
    Current state                                                :        Normal
    Overload mode count                                          :             0
    Self-protection mode count                                   :             0
    Last detected Overload state                                 :   0000:00:00:00:00:00

    Last detected Out-of Overload state                          :   0000:00:00:00:00:00

    Last reported (system) Overload Factor
      CPU                                                        :             0
      Memory                                                     :             0
      Session Capacity                                           :             0
    Last reported (VPP CPU) Overload Factor                      :             0
    Last reported Reduction Metric                               :             0

   Sessions Stats:
      Sx-Sessions removed in Self-protection mode                :             0
      Non-Priority Sx Messages:
        Sx-Establishment reqs rcvd in Overload mode              :             0
        Sx-Establishment reqs rejected in Self-protection mode   :             0
        Sx-Modification reqs rejected in Self-protection mode    :             0
      Unclassified Sx Messages:
        Sx-Establishment reqs rcvd in Overload mode              :             0
        Sx-Establishment reqs rejected in Self-protection mode   :             0
        Sx-Modification reqs rejected in Self-protection mode    :             0
      IMS Sx Messages:
        Sx-Establishment reqs rcvd in Overload mode              :             0
        Sx-Establishment reqs rejected in Self-protection mode   :             0
        Sx-Modification reqs rejected in Self-protection mode    :             0
        Sx-Establishment reqs allowed in Self-protection mode    :             0
        Sx-Modification reqs allowed in Self-protection mode     :             0
      WPS/Emergency Sx Messages:
        Sx-Establishment reqs rcvd in Overload mode              :             0
        Sx-Establishment reqs rejected in Self-protection mode   :             0
        Sx-Modification reqs rejected in Self-protection mode    :             0
```

```
      Sx-Establishment reqs allowed in Self-protection mode         :              0
      Sx-Modification reqs allowed in Self-protection mode          :              0

  Data Stats:
    Overload Mode:
      WPS/Emergency new flows allowed                               :              0
      WPS/Emergency new flows dropped                               :              0
      Non WPS/Emergency new flows allowed                           :              0
      Non WPS/Emergency new flows dropped                           :              0
      Unclassified new flows dropped                                :              0
      Total Packets dropped                                         :              0
    Self-Protection Mode:
      WPS/Emergency new flows allowed                               :              0
      WPS/Emergency new flows dropped                               :              0
      Non WPS/Emergency new flows allowed                           :              0
      Non WPS/Emergency new flows dropped                           :              0
      Unclassified new flows dropped                                :              0
      Total Packets dropped                                         :              0

UPF Load Stats:
  Current (system) Load Factor
    CPU                                                             :              0
    Memory                                                          :              0
    Session Capacity                                                :              0
  Last reported sequence number                                     :              0
```

show user-plane-service statistics all

The show command **show user-plane-service statistics all** displays the load and overload stats for UPF.

```
UPF Overload Stats:
  Current state                                                     :         Normal
  Overload mode count                                               :              0
  Self-protection mode count                                        :              0
  Last detected Overload state                                    :   0000:00:00:00:00:00

  Last detected Out-of Overload state                             :   0000:00:00:00:00:00

  Last reported (system) Overload Factor
    CPU                                                             :              0
    Memory                                                          :              0
    Session Capacity                                                :              0
  Last reported (VPP CPU) Overload Factor                           :              0
  Last reported Reduction Metric                                    :              0

 Sessions Stats:
    Sx-Sessions removed in Self-protection mode                     :              0
    Non-Priority Sx Messages:
      Sx-Establishment reqs rcvd in Overload mode                   :              0
      Sx-Establishment reqs rejected in Self-protection mode        :              0
      Sx-Modification reqs rejected in Self-protection mode         :              0
    Unclassified Sx Messages:
      Sx-Establishment reqs rcvd in Overload mode                   :              0
      Sx-Establishment reqs rejected in Self-protection mode        :              0
      Sx-Modification reqs rejected in Self-protection mode         :              0
    IMS Sx Messages:
      Sx-Establishment reqs rcvd in Overload mode                   :              0
      Sx-Establishment reqs rejected in Self-protection mode        :              0
      Sx-Modification reqs rejected in Self-protection mode         :              0
      Sx-Establishment reqs allowed in Self-protection mode         :              0
      Sx-Modification reqs allowed in Self-protection mode          :              0
    WPS/Emergency Sx Messages:
      Sx-Establishment reqs rcvd in Overload mode                   :              0
      Sx-Establishment reqs rejected in Self-protection mode        :              0
```

```
        Sx-Modification reqs rejected in Self-protection mode         :            0
        Sx-Establishment reqs allowed in Self-protection mode         :            0
        Sx-Modification reqs allowed in Self-protection mode          :            0

    Data Stats:
      Overload Mode:
        WPS/Emergency new flows allowed                               :            0
        WPS/Emergency new flows dropped                               :            0
        Non WPS/Emergency new flows allowed                           :            0
        Non WPS/Emergency new flows dropped                           :            0
        Unclassified new flows dropped                                :            0
        Total Packets dropped                                         :            0
      Self-Protection Mode:
        WPS/Emergency new flows allowed                               :            0
        WPS/Emergency new flows dropped                               :            0
        Non WPS/Emergency new flows allowed                           :            0
        Non WPS/Emergency new flows dropped                           :            0
        Unclassified new flows dropped                                :            0
        Total Packets dropped                                         :            0

  UPF Load Stats:
    Current (system) Load Factor
      CPU                                                             :            0
      Memory                                                          :            0
      Session Capacity                                                :            0
    Last reported sequence number                                     :            0
```

show user-plane-service name

The show command **show user-plane-service name** *service_name* displays the load and overload control profile information along with the other information for a specific user plane service.

```
show user-plane-service name UPlane1

Service name                        : UPlane1
  Service-Id                        : 6
  Context                           : EPC
  Status                            : STARTED
  PGW Ingress GTPU Service          : pgwIngress
  UPF Ingress N3 GTPU Service       : Not defined
  UPF Ingress N9 GTPU Service       : Not defined
  UPF Ingress S5U GTPU Service      : Not defined
  UPF Ingress S8U GTPU Service      : Not defined
  UPF Ingress N9_S5U_S8U GTPU Service: Not defined
  UPF Egress GTPU Service           : Not defined
  SGW Ingress GTPU Service          : SGW1EGTPINGTPUIN
  SGW Egress GTPU Service           : SGW1EGTPOUTGTPUOUT
  Control Plane Tunnel GTPU Service : UP1SXU
  Sx Service                        : sx1
  Control Plane Group               : cp1
  Userplane Load Control Profile    : Not defined
  Userplane Overload Control Profile: Not defined
  Fast-Path service                 : Disabled
  NRF Management Service            : Not defined
  Nf-Instance-Id                    : Not defined
  Session Priority Profile          : Not defined
  Ise Server Profile                : Not defined
  UPF Load Control Profile          : ULCI1
  UPF Overload Control Profile      : UOCI1
```

show user-plane-service all

The show command **show user-plane-service all** displays the load and overload control profile information along with other information for a user plane service.

```
show user-plane-service all

Service name                        : UPlane1
  Service-Id                        : 6
  Context                           : EPC
  Status                            : STARTED
  PGW Ingress GTPU Service          : pgwIngress
  UPF Ingress N3 GTPU Service       : Not defined
  UPF Ingress N9 GTPU Service       : Not defined
  UPF Ingress S5U GTPU Service      : Not defined
  UPF Ingress S8U GTPU Service      : Not defined
  UPF Ingress N9_S5U_S8U GTPU Service: Not defined
  UPF Egress GTPU Service           : Not defined
  SGW Ingress GTPU Service          : SGW1EGTPINGTPUIN
  SGW Egress GTPU Service           : SGW1EGTPOUTGTPUOUT
  Control Plane Tunnel GTPU Service : UP1SXU
  Sx Service                        : sx1
  Control Plane Group               : cp1
  Userplane Load Control Profile    : Not defined
  Userplane Overload Control Profile: Not defined
  Fast-Path service                 : Disabled
  NRF Management Service            : Not defined
  Nf-Instance-Id                    : Not defined
  Session Priority Profile          : Not defined
  Ise Server Profile                : Not defined
  UPF Load Control Profile          : ULCI1
  UPF Overload Control Profile      : UOCI1
```

show sx-service statistics all

The show command **show sx-service statistics all** displays the session management messages terminated in self-protection state.

```
show sx-service statistics all

Session Management Messages:
….
Session Report Request:
 Total TX:                          1   Total RX:                          0
 Initial TX:                        1   Initial RX:                        0
 Retrans TX:                        0   Retrans RX:                        0
 Discarded:                         0   No Rsp RX:                         0
 Causes:
  User Initiated Session Deletion from UP:                                 0
  Association Release initiated by UP:                                    0
  Recovery Failure:                                                        0
  IP Source Violation:                                                     0
  Self Protection Termination:                                             1
```

Bulkstats

Following bulkstats are supported as part of this feature:

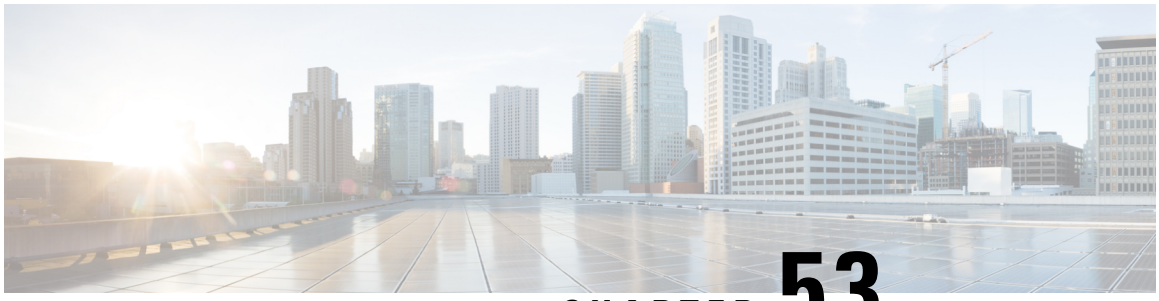
Variable Name	Data Type	Key	Counter Type
%upf-num-overload-state-reached%	Int32	0	Counter
%upf-num-self-protection-reached%	Int32	0	Counter
%upf-last-in-overload-state-timestamp%	Int32	0	Info
%upf-last-out-overload-state-timestamp%	Int32	0	Info
%upf-overload-reduction-metric%	Int32	0	Info
%upf-overload-factor-system-cpu%	Int32	0	Info
%upf-overload-factor-system-mem%	Int32	0	Info
%upf-overload-factor-system-sesscnt%	Int32	0	Info
%upf-overload-factor-vpp-cpu%	Int32	0	Info
%upf-load-factor-system-cpu%	Int32	0	Info
%upf-load-factor-system-mem%	Int32	0	Info
%upf-load-factor-system-sesscnt%	Int32	0	Info
%upf-load-factor-seq-num%	Int32	0	Info
%upf-num-sessions-removed-during-self-protection-mode%	Int32	0	Info
%upf-num-non-priority-session-estab-req-received-in-overload%	Int32	0	Counter
%upf-num-non-priority-session-estab-req-reject-in-sp%	Int32	0	Counter
%upf-num-non-priority-session-mod-req-reject-in-sp%	Int32	0	Counter
%upf-num-unclassified-session-estab-req-received-in-overload%	Int32	0	Counter
%upf-num-unclassified-session-estab-req-reject-in-sp%	Int32	0	Counter
%upf-num-unclassified-session-mod-req-reject-in-sp%	Int32	0	Counter
%upf-num-ims-session-estab-req-received-in-overload%	Int32	0	Counter
%upf-num-ims-session-estab-req-reject-in-sp%	Int32	0	Counter
%upf-num-ims-session-mod-req-reject-in-sp%	Int32	0	Counter
%upf-num-ims-session-estab-req-allowed-in-sp%	Int32	0	Counter
%upf-num-ims-session-mod-req-allowed-in-sp%	Int32	0	Counter
%upf-num-wps-emer-session-estab-req-received-in-overload%	Int32	0	Counter
%upf-num-wps-emer-session-estab-req-reject-in-sp%	Int32	0	Counter

Variable Name	Data Type	Key	Counter Type
%upf-num-wps-emer-session-mod-req-reject-in-sp%	Int32	0	Counter
%upf-num-wps-emer-session-estab-req-allowed-in-sp%	Int32	0	Counter
%upf-num-wps-emer-session-mod-req-allowed-in-sp%	Int32	0	Counter
%upf-num-wps-emer-new-flows-allowed-in-overload%	Int32	0	Counter
%upf-num-wps-emer-new-flows-dropped-in-overload%	Int32	0	Counter
%upf-num-non-wps-emer-new-flows-allowed-in-overload%	Int32	0	Counter
%upf-num-non-wps-emer-new-flows-dropped-in-overload%	Int32	0	Counter
%upf-num-unclassified-new-flows-dropped-in-overload%	Int32	0	Counter
%upf-num-packets-dropped-in-overload%	Int32	0	Counter
%upf-num-wps-emer-new-flows-allowed-in-self-protection%	Int32	0	Counter
%upf-num-wps-emer-new-flows-dropped-in-self-protection%	Int32	0	Counter
%upf-num-non-wps-emer-new-flows-allowed-in-self-protection%	Int32	0	Counter
%upf-num-non-wps-emer-new-flows-dropped-in-self-protection%	Int32	0	Counter
%upf-num-unclassified-new-flows-dropped-in-self-protection%	Int32	0	Counter
%upf-num-packets-dropped-in-self-protection%	Int32	0	Counter

SNMP Traps

The following SNMP Traps are added in support of this feature:

- **UPFOverloadState**: This trap is generated when the system enters the overload state.
- **UPFOverloadStateClear**: This trap is generated when the system is out of the overload state.
- **UPFSelfProtectionMode**: This trap is generated when the system enters the self-protection mode.
- **UPFSelfProtectionModeClear**: This trap is generated when the system is out of the self-protection mode.



CHAPTER 53

UPF Usage Monitoring over PCF

- Feature Summary and Revision History, on page 583
- Feature Description, on page 584

Feature Summary and Revision History

Summary Data

Table 124: Summary Data

Applicable Product(s) or Functional Area	5G-UPF
Applicable Platform(s)	VPC-SI SMI
Feature Default Setting	Enabled - Always-on
Related Changes in This Release:	Not Applicable
Related Documentation	• *UCC 5G SMF Configuration and Administration Guide* • *UCC 5G UPF Configuration and Administration Guide*

Revision History

Revision Details	Release
First introduced.	2022.01.0

Feature Description

UPF supports usage monitoring control over the existing N4 interface to report usage thresholds that are provided from PCF over the N7 interface through SMF for both 4G and 5G sessions. UPF reports the usage threshold breach to SMF through Session Report Request and SMF sends the data to PCF. UPF supports the modification of usage monitoring parameters, such as Total Volume, Uplink Volume, or Downlink Volume thresholds and the disabling of usage monitoring based on non-reception of usage monitoring threshold or related triggers from PCF.

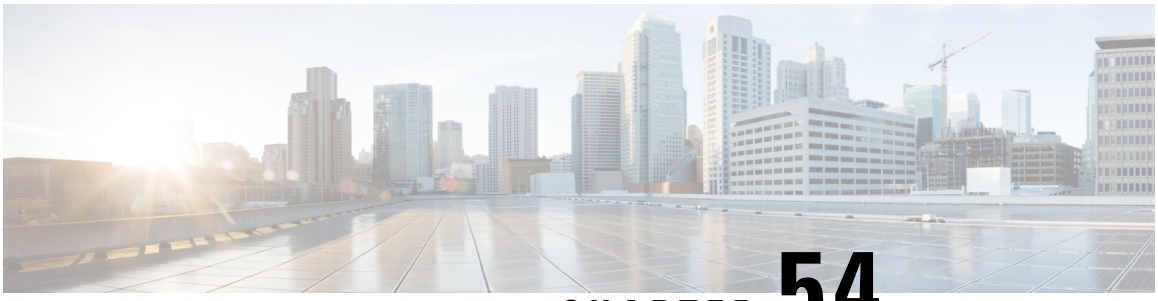
Usage Reporting

UPF measures the volume and the time usage of all traffic for the PDU session or the corresponding service data flows. UPF sends the accumulated usage report in either the PFCP Session Report Request or the PFCP Session Modification Response to SMF. SMF includes one or multiple accumulated usage reports in the "accuUsageReports" attribute towards PCF.



Note The *Usage Monitoring over PCF* feature is enabled from SMF.

NOTE: To know more about how SMF handles this functionality, refer to the *Usage Monitoring over PCF* section in the *Policy and User Plane Management* chapter of *UCC 5G SMF Configuration and Administration Guide*.



CHAPTER 54

Variable Relinquishing of CPU Cycles

Table 125: Feature History

Feature Name	Release Information	Description
Power Management using VRCC	2024.04.0	The Variable Relinquishing of CPU Cycles (VRCC) feature in 5G UPF introduces a dynamic CPU management system to reduce power consumption in packet processing systems based on traffic demands. This feature allows the CPU to reduce its clock frequency during low traffic periods and engage the power saving mode. This feature requires a valid license and is supported on VPP instances of VPC-SI for 5G UPF only. **Command Introduced**: **require power-saving vpp** in Global Configuration mode **Default Setting**: Disabled – Configuration Required to Enable

- VRCC Overview, on page 585
- How VRCC Works, on page 586
- Configure VRCC, on page 587

VRCC Overview

About VPP

The 5G UPF network function utilizes a third-party Vector Packet Processing (VPP) system to process packets in batches called vectors. Using VPP, UPF achieves ultra-fast packet forwarding and integration with IP services and third-party applications for compatibility with all user plane functionality.

About VRCC

The Variable Relinquishing of CPU Cycles (VRCC) feature in 5G UPF allows the system to relinquish unused processing power during low-traffic periods. This feature allows UPF to dynamically manage CPU cycles based on packet arrival rates. It ensures optimal resource utilization and power efficency without impacting traffic processing or user experience.

Supported Platforms for VRCC

The VRCC feature runs on VPP instances of VPC-SI for 5G UPF only.

License Requirements for VRCC

You need a license to enable or disable the VRCC feature.

Contact your Cisco account representative for more information.

Benefits of VRCC

The VRCC feature provides these benefits:

- Energy Efficiency—Reduces power consumption during low traffic periods with measurable energy savings.
- Environmental Impact—Reduces the carbon footprint for network operations.
- Cost Savings—Decreases costs associated with power and cooling.
- Dynamic Adjustment—Adjusts CPU performance automatically based on real-time traffic demands.

How VRCC Works

VRCC is self-contained within each VPP worker thread and integrates with the packet processing engine's scheduler. VPP runs poll mode drivers (PMD) to fully uitilize the CPU regardless of the interface load.

VRCC computes a moving average of VPP utilization to gauge the traffic load and periodically relinquishes CPU cycles for short periods. This mechanism enables the system to balance throughput, latency response, system stability, and power management.

Detect High and Low Traffic

The main loop scheduler of each VPP thread computes a sleep length. The sleep length is inversely proportional to a moving average of the incoming vector size. The vector size indicates the high and low traffic to detect the system performance.

- High Traffic—The moving average is large with short or no sleep periods. This state keeps the CPU active.
- Low Traffic—The moving average is large with longer sleep periods. This state reduces CPU activity.

Engage Power Saving Mode

VRCC applies short sleep periods within the VPP's processing loop. The CPU pauses processing temporarily during these periods. By applying a voluntary short sleep, the VPP worker threads allow the CPU to change its clock frequency dynamically in real-time. It triggers the CPU hardware indirectly to engage the power saving modes in response to short bursts of sleep.

Effects on Network Latency

In real-world scenarios, the expected round-trip latency is typically 40 milliseconds or more, according to Ookla benchmarking data.

When the VRCC feature is enabled, there is usually a slight increase in latency, approximately 0.3 to 0.4 milliseconds, based on lab tests. This increase accounts for only about 1% to 2% of the total round-trip latency expected in real-world conditions.

This additional latency occurs because the VPP must resume from a VRCC-induced sleep to handle incoming traffic bursts. If VPP remains active, these sleep cycles don't happen, and latency remains the same. However, if a burst arrives while VPP is idle, there will be a brief delay as the system returns to full responsiveness.

Configure VRCC

Use this task to engage the CPU hardware in VPP power savings mode.

Before you begin

- Ensure that you are running the correct call model based on your deployment.
- You require a valid license to enable the VRCC feature.

Procedure

Step 1 Enable the power savings mode using the **require power-saving vpp** CLI command. This command is enforced during runtime and does not require reload.

Example:

```
[local]host_name# configure
[local]host_name(config)# require power-saving vpp
[local]host_name(config)# exit
```

Use the **no require power-saving vpp** CLI command to disable the power savings mode.

Step 2 Verify whether VRCC is enabled using the **show npumgr vpp vrcc** command.

Note
You can configure the **show npumgr vpp vrcc** CLI only after entering the hidden mode.

Example:

```
[local]host_name# show npumgr vpp vrcc
VPP VRCC : Enabled
```

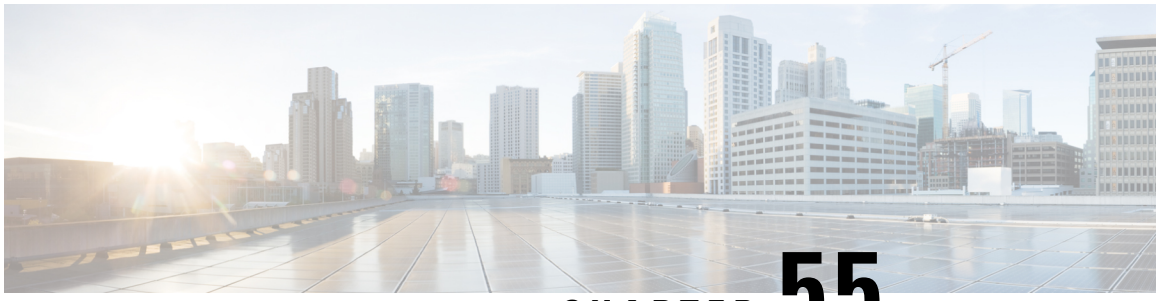
Step 3 On your host CIMC, check the Power Management or Power Monitoring tab to view the graph for power measurement. See the component CPU in the graph to view the CPU power domain details.

You can access power monitoring from the CIMC CLI with the following commands:

Example:

```
[server]# scope chassis
[server]# show power-monitoring detail
```

Step 4 Check the system KPIs and monitor the system for any crash or abnormal behaviour. Also, perform your regular resiliency or redundancy tasks.



CHAPTER 55

Virtual Routing and Forwarding

- Feature Summary and Revision History, on page 589
- Revision History, on page 589
- Feature Description, on page 589
- Configuring VRF, on page 592
- VLAN Segmentation Using VRF, on page 595
- Monitoring and Troubleshooting, on page 596

Feature Summary and Revision History

Revision History

Table 126: Revision History

Revision Details	Release
UPF supports up to 200 VRFs for private APN/DNN.	2023.02.0
UPF supports up to 129 VRFs for private APN/DNN.	2022.04.0
Support is added for the following functionality: • Overlapping IP Pools • Removal of mandatory VRF ordering between SMF and UPF	2021.01.0
First introduced.	2020.02.0

Feature Description

Virtual Routing and Forwarding (VRF) is a technology that allows multiple instances of a routing table to coexist within the same router at the same time. As the routing instances are independent, VRF uses the same or overlapping IP addresses without conflicting with each other.

Note In 2023.02.0 and later releases, UPF supports up to 200 VRFs for private APN or DNN.

In releases prior to 2023.02.0, UPF supported up to 129 VRFs for private APN or DNN.

In UPF, this feature enables association of IP address pools with VRF. The chunks from this pool are allocated to the UPFs that are configured to use these pools. VRF-associated pools in UPF can be either Static or Private type.

When UPF comes up for registration, the chunks in the PRIVATE VRF pool are allocated similar to the normal private pools. For a Static VRF pool, SMF does chunk allocation to UPF during configuration. An Sx-Route-Update message is sent for pre-allocated static chunks during UPF registration.

Overlapping IP Pool

Overlapping pools share and use an IP address range. Overlapping pools can either be of Static or Private type. Public pools cannot be configured as overlapping pools. Each overlapping pool is part of a different VRF (routing domain) and pool-group. Since an APN can use only one pool-group, overlapping pools are part of different APNs.

Without this functionality, overlapping pools are configured at SMF. However, chunks from two overlapping pools cannot be sent to the same UPF. That is, the UPF cannot handle chunks from two different overlapping pools. Same number of UPFs and overlapping pools are required for sharing the same IP address range.

With this functionality, UPF handles chunks from two different overlapping pools. So, a single UPF can handle any number of overlapping pools sharing the same IP range.

The functionality of overlapping pools in the same UPF includes:

- When a chunk from a particular pool is installed on UPF, its corresponding vrf-name is sent along with the chunk.
- The UPFs are VRF-aware of the chunks and install chunks on the corresponding VRFs. The chunk database is populated under VRFs.
- During call allocation, release, recovery, or any communication toward VPNMgr, the corresponding SessMgr at UPF includes vrf-id. This enables VPNMgr to select the correct chunk for that IP under the provided vrf-id for processing.

UE IP VRF is a custom IE that encapsulates the VRF name of N4 SESSION ESTABLISHMENT REQUEST message.

SessMgr in UPF converts the received vrf- name into UPF VRF CONTEXT ID and passes it on to UPFs VPNMgr in IP allocation request. This UPF VRF CONTEXT ID is also used when release request is sent to UPFs VPNMgr. UPFs VPNMgr is made aware of the VRF to which that chunk belongs to by sending vrf-name in each chunk-related communication between SMF and UPF. This enables UPFs VPNMgr to create database of chunks under each VRF enabling support of overlapping pools in the same UPF.

UE IP VRF Information Element

The following is the IE format of the private UE IP VRF.

Table 127: UE IP VRF Format

	Bits							
Octets	8	7	6	5	4	3	2	1
1 to 2	Type = 242 (decimal)							
3 to 4	Length = n							
5	Spare					Identical VRF flag	IPv6 VRF Valid	IPv4 VRF Valid
m to m+1	VRF-1 Name Length = p							
m+1 to m+1+p	VRF-1 Name							
n to n+1	VRF-2 Name Length = q							
m+1 to m+1+q	VRF-2 Name							

The following table shows the possible values of the "UE IP VRF" fields.

Cases	UE IP VRF	Value (binary)		
		Bit 3	Bit 2	Bit 1
1	None of the IPv4 and IPv6 UE IP addresses are associated to VRF.	0	0	0
2	Only IPv4 UE IP address is associated to a VRF.	0	0	1
3	Only IPv6 UE IP address is associated to VRF.	0	1	0
4	Both IPv4 and IPv6 UE IP addresses are associated to different VRFs.	0	1	1
5	Both IPv4 and IPv6 UE IP addresses are associated to a common VRF.	1	1	1

VRF Name as Identifier

The communication between SMF and UPF, related to VRF, was done through vrf-id. This required the operator to have all VRFs configured in both SMF and UPF, and also in the same order.

With this feature, vrf-name is used as identifier in all the communication between SMF and UPF related to VRFs. This feature eliminates the configuration of all VRFs in UPF. Operator can configure VRFs in different order at SMF and UPF, and can identify the VRF with the same vrf-name in both the nodes.

Limitations and Restrictions

The following are the known limitations and restrictions in UPF:

- VRF limit per UPF is 205.

- UPF supports only VRF-based overlapping pools. UPF does not support overlapping pools such as NH-based and VLAN-based pools.
- UPF does not permit PDN Type IPv4v6-based call on static IP pools with multiple UPFs in the same UPF group.
- UPF does not support dynamic update of VRF.

Configuring VRF

Use the following steps to configure VRF support in UPF.

At SMF:

1. Create the APN or DNN profile.
2. Create overlapping IP pools and associate the respective APN or DNN and VRF at context-level.
3. Associate APN or DNN to the UPF profile.

The following is an example of the SMF configuration:

```
profile dnn intershat1
.
.
.
   upf apn mpls1.com
exit
profile dnn intershat2
.
.
.
   upf apn mpls2.com
exit
profile network-element upf upf1
.
.
.
   dnn-list [ intershat1 intershat2 ]
exit
profile network-element upf upf2
.
.
.
   dnn-list [ intershat1 intershat2 ]
exit
ipam
   source local
   address-pool pool-intershat1
      vrf-name mpls-vrf-1@isp
      tags
         dnn intershat1
      exit
      ipv4
         address-range 209.165.201.25 255.255.255.224
      exit
   exit
   address-pool pool-intershat2
      vrf-name mpls-vrf-2@isp
      tags
```

```
        dnn intershat2
      exit
      ipv4
         address-range 209.165.201.25 255.255.255.224
      exit
   exit
exit
```

At UPF:

It is recommended to configure VRF in UPF before a chunk is pushed from SMF. Else, it leads to failure of the complete IP pool transaction (including chunks that donot belong to the VRF). SMF retries the attempt after some time.

1. Create VRFs in the same order as it is created in the CP.
2. Create at least one interface per VRF to carry traffic for the chunks allocated in that VRF.

The following is an example of the UPF configurations:

UPF 1:

```
config
  context EPC2
    sx-service sx
      instance-type userplane
      bind ipv4-address 209.165.201.11 ipv6-address bbbb:aaaa::4
    exit
    user-plane-service up
      associate gtpu-service pgw-gtpu pgw-ingress
      associate gtpu-service sgw-ingress-gtpu sgw-ingress
      associate gtpu-service sgw-engress-gtpu sgw-egress
      associate gtpu-service saegw-sxu cp-tunnel
      associate sx-service sx
      associate fast-path service
      associate control-plane-group g1
    exit

  context isp
    ip vrf mpls-vrf-1
    #exit
    ip vrf mpls-vrf-2
    #exit
    ip vrf mpls-vrf-1
        route-distinguisher 61601 11100001
        route-target export 61601 11100001
        route-target import 61606 11100001
        route-target import 65200 11100001
    #exit
    address-family ipv4 vrf mpls-vrf-1
        redistribute connected
        redistribute static
    #exit
    address-family ipv6 vrf mpls-vrf-1
        redistribute connected
        redistribute static
    #exit
    ip vrf mpls-vrf-2
        route-distinguisher 61601 11100002
        route-target export 61601 11100002
        route-target import 61606 11100002
        route-target import 65200 11100002
    #exit
    address-family ipv4 vrf mpls-vrf-2
```

```
        redistribute connected
        redistribute static
    #exit
    address-family ipv6 vrf mpls-vrf-2
        redistribute connected
        redistribute static
    #exit
    #exit
apn mpls1.com
      pdp-type ipv4 ipv6
      bearer-control-mode mixed
      selection-mode sent-by-ms
      ip context-name isp
    exit
exit
control-plane-group g1
    peer-node-id ipv4-address 209.165.201.15
  #exit
  user-plane-group default
```

UPF 2:

```
config
  context EPC2
    sx-service sx
      instance-type userplane
      bind ipv4-address 209.165.201.12 ipv6-address bbbb:aaaa::5
    exit
    user-plane-service up
      associate gtpu-service pgw-gtpu pgw-ingress
      associate gtpu-service sgw-ingress-gtpu sgw-ingress
      associate gtpu-service sgw-engress-gtpu sgw-egress
      associate gtpu-service saegw-sxu cp-tunnel
      associate sx-service sx
      associate fast-path service
      associate control-plane-group g1
    exit
exit

  context isp
    ip vrf mpls-vrf-1
    #exit
    ip vrf mpls-vrf-2
    #exit
    apn mpls2.com
      pdp-type ipv4 ipv6
      bearer-control-mode mixed
      selection-mode sent-by-ms
      ip context-name isp
    exit
exit

control-plane-group g1
    peer-node-id ipv4-address 209.165.201.15
  #exit
  user-plane-group default
```

VLAN Segmentation Using VRF

Table 128: Feature History

Feature Name	Release Information	Description
Enabling VLAN Segmentation of Traffic Using Virtual Routing and Forwarding	2024.01	UPF allows VLAN segmentation using VRF. A VRF is configured for each VLAN and each VRF has a default route through which the uplink packets are sent. **Default Setting**: Disabled-Configuration required to enable.

Feature Description

UPF supports segmenting the enterprise traffic using Virtual LAN (VLAN). VLANs enable a single physical LAN segment to be further segmented into smaller virtual networks so that the groups of ports are isolated from one another, as if they were on physically different segments. In VLAN segmentation, the UEs belonging to different classes are assigned to different VLANs for traffic isolation.

VLAN Segmentation can be achieved using two methods:

1. Next Hop Forwarding
2. Virtual Routing and Forwarding

UPF supports segmenting VLAN traffic by configuring VRF. It is a layer-3 virtualization that consists of configuring a VRF with a defaut route per VLAN. The VRF default route is used for sending packets.

VLAN segmentation is enabled by logically mapping the VLAN tags with the IP pools at the SMF. The UPF is informed about the mapping configuration based on which UPF forwards the uplink packets on the N6 interface through the default route configured for a VRF.

Enabling VLAN Segmentation Using VRF

To enable VLAN segmentation using VRF on UPF, use following configuration:

```
config
  port ethernet slot/port
    no shutdown
    vlan vlan_tag_ID
    no shutdown
    bind interface interface_name context_name
    end
    context context_name
      ip vrf vrf_name
      { ip | ipv6 } address address subnetmask [ secondary ]
      end
```

NOTES:

- **ip vrf** *vrf_name*—Creates a VRF in the router and assigns a VRF-ID.

Configuration Example

Following is a sample configuration for enabling VLAN segmentation using VRF:

```
config
    port ethernet 1/10
       no shutdown
       vlan 400
         no shutdown
         bind interface N6_interface1 ISP
       #exit
       vlan 401
         no shutdown
         bind interface N6_interface2 ISP
       #exit
 context ISP
    ip vrf vrf1
    #exit
    ip vrf vrf2
    #exit
 interface N6_interface1
      ip vrf forwarding vrf1
      ip address 209.165.200.225 209.165.200.254
      ip address 209.165.201.1 209.165.200.254 secondary
      ipv6 address 2001:DB8::1/32 secondary
 interface N6_interface2
      ip vrf forwarding vrf2
      ip address 209.165.202.129 209.165.200.254
      ip address 209.165.201.30 209.165.200.254 secondary
      ipv6 address 2001:DB8:1::1/32 secondary
```

Monitoring and Troubleshooting

This section provides information regarding the CLI commands available for monitoring and troubleshooting the feature.

Show Commands and Outputs

This section provides information regarding show commands and their outputs in support of this feature.

show ip chunks

The output of this CLI command displays all chunks in that context.

With the Overlapping IP Pools functionality, the **show ip chunks vrf** *vrf_name* CLI command displays only chunks under that VRF.

- chunk-id
- chunk-size
- vrf-name

- start-addr
- end-addr
- used-addrs
- Peer Address

Following is a sample output:

```
========================================
VRF Name: MPN00001
=============================================================================================================
|------------------------------|----------|------------|--------------|--------------|------------|
|          Peer Address          | chunk-id | chunk-size |   start-addr  |    end-addr   |
 used-addrs |
|------------------------------|----------|------------|--------------|--------------|------------|
|                     192.10.25.23|1074790405|        8192|    36.40.128.0|  36.40.159.255|
           0|
|                     192.10.25.23|1074790406|        8192|    36.40.160.0|  36.40.191.255|
           0|
|------------------------------|----------|------------|--------------|--------------|------------|
```

show ipv6 chunks

The output of this CLI command displays all chunks in that context.

With the Overlapping IP Pools functionality, the output of the **show ipv6 chunks vrf** *vrf_name* CLI command displays only chunks under that VRF.

- chunk-id
- chunk-size
- vrf-name
- start-prefix
- end-prefix
- used-prefixes
- Peer Address

show ip bgp vpnv4

The output of this CLI command displays all the VPN routing information.

With the Overlapping IP Pools functionality, the **show ip bgp vpnv4 vrf** *vrf_name* CLI command displays the information under that VRF.

Following is a sample output:

```
VPNv4 Routing Table:
BGP table version is 1, local router ID is 172.31.35.36
Status Codes: s suppressed, d damped, h history, * valid, > best, i - internal, S stale, m
 Multipath
Origin Codes: i - IGP, e - EGP, ? - incomplete

Network              Next Hop             Metric LocPrf Weight Path
*>   0.0.0.0/0            172.31.35.29         0             0      65200 ?
```

```
*>   2.2.2.101/32         172.31.35.29         0             0      65200 ?
*>   2.2.3.2/32           0.0.0.0              0             32768  ?
*>   36.40.0.0/19         0.0.0.0              0             32768  ?
*>   36.40.32.0/19        0.0.0.0              0             32768  ?
*>   36.40.192.0/19       0.0.0.0              0             32768  ?

Total number of prefixes 6
```

show ip bgp vpnv6

The output of this CLI command displays all the VPN routing information.

With the Overlapping IP Pools functionality, the **show ip bgp vpnv6 vrf** *vrf_name* CLI command displays the information under that VRF.

Following is a sample output:

```
VPNv6 Routing Table:
BGP table version is 1, local router ID is 172.31.35.36
Status Codes: s suppressed, d damped, h history, * valid, > best, i - internal, S stale, m
 Multipath
Origin Codes: i - IGP, e - EGP, ? - incomplete

 Network              Next Hop             Metric LocPrf Weight Path
*>   0.0.0.0/0            172.31.35.29         0             0      65200 ?
*>   2036:adb0:40         172.31.35.29/51         0             32768  ?
*>   2036:adb0:40:2000    172.31.35.29/51           0             32768  ?
*>   2036:adb0:40:c000    172.31.35.29/51            0             32768  ?

Total number of prefixes 4
```

show mpls ilm

The output of this CLI command displays the MPLS ILM table with FEC information.

With the Overlapping IP Pools functionality, the **show mpls ilm** *fec verbose* and **show mpls ilm** *fec summary* CLI commands display the information under the VRF.

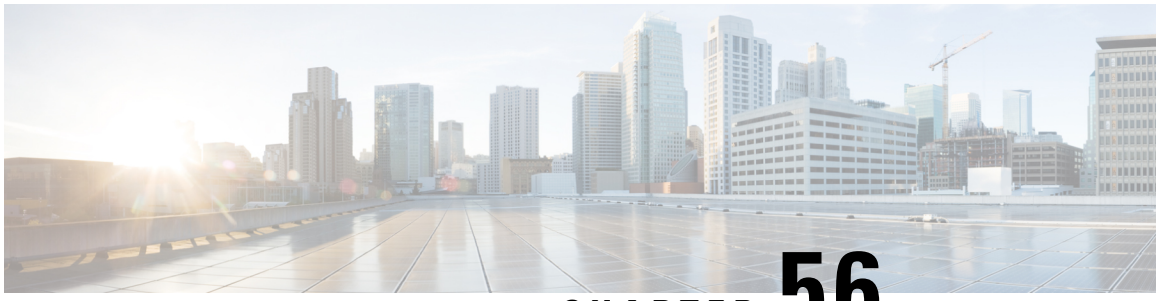
The following is a sample output for the **show mpls ilm** *fec summary* CLI command to display the count.

```
Total ILM entries: 406
```

The following is a sample output for the **show mpls ilm** *fec Verbose* CLI command to display a detailed MPLS ILM table.

```
In-segment entry with in label: 832, id: 99, in label-space: 0, row status: Active
Owner: BGP, # of pops: 1 XC Index:708
FEC: 36.40.0.0/19
FEC: 36.40.32.0/19
FEC: 36.40.192.0/19

In-segment entry with in label: 833, id: 100, in label-space: 0, row status: Active
Owner: BGP, # of pops: 1 XC Index:709
FEC: 2037:adb0:201::/51
FEC: 2037:adb0:201::/51
FEC: 2037:adb0:201::/51
```



CHAPTER 56

Voice over New Radio

This chapter covers the following topics:

- Feature Summary and Revision History, on page 599
- Feature Description, on page 600

Feature Summary and Revision History

Summary Data

Table 129: Summary Data

Applicable Product(s) or Functional Area	5G-UPF
Applicable Platform(s)	VPC-SI SMI
Feature Default Setting	Enabled – Always-on
Related Changes in this Release	Not Applicable
Related Documentation	Not Applicable

Revision History

Table 130: Revision History

Revision Details	Release
First Introduced.	2020.02.0

Feature Description

The UPF supports Voice over New Radio (VoNR) with the existing Session Establishment and Modification procedures. In these procedures, the SMF creates the PDR for 5QI=5 Non-GBR flow for IMS signaling and PDR for 5QI=1 GBR flow for voice traffic. The UPF does not require any special handling to support mobile-originated or mobile-terminated call flows.

How it Works

The following are the steps in the call flow in which the PDRs are created with 5QI value 5 for IMS signaling and 5QI value 1 or Voice Traffic.

VoNR Call Flow for UPF

This section describes the steps for VoNR session and respective PDR Creation on UPF.

Figure 34: VoNR Call Flow

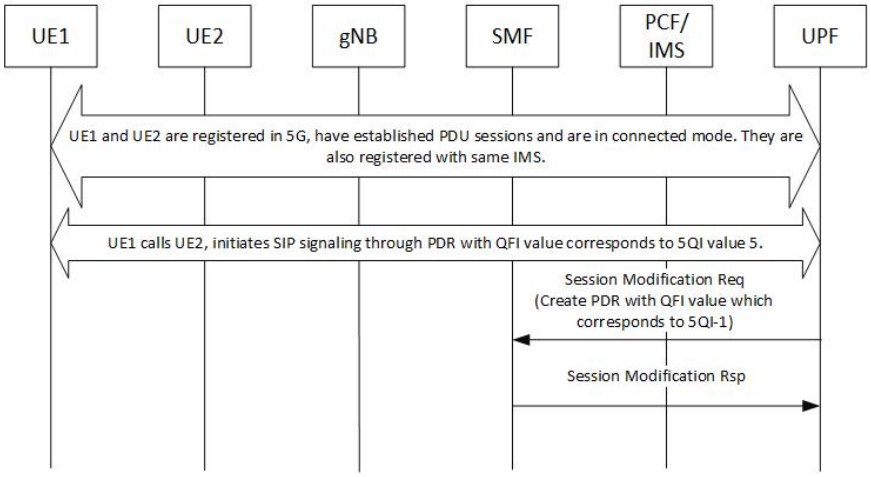
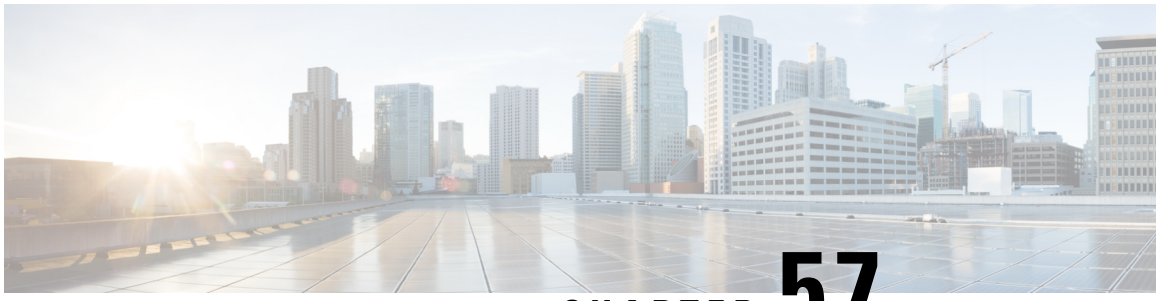


Table 131: VoNR Call Flow

Step	Description
1	UE1 and UE2 are registered on 5G network. They establish the IMS PDU session and both are registered with same IMS. Both UEs are in connected mode.
2	For IMS signaling, a non-GBR QoS PDRs (UL and DL) is created by SMF which has the QFI value that corresponds to 5QI value 5.
3	Similarly, for Conversational Voice traffic, the PDRs for GBR QoS flow is created with the QFI value that corresponds to 5QI value 1.
4	The QFI to 5QI mapping is maintained by SMF, hence the QFI does not have the values same as 5QI.

Step	Description
5	The above steps are valid for both Mobile Originated (MO) or Mobile Terminate (MT) call flows.
6	Refer to 3GPP TS 23.501, Section 5.7.4 for other types of 5QI mappings for GBR or Non-GBR flows.



CHAPTER 57

WPS Prioritization

Table 132: Feature History

Feature Name	Release Information	Description
WPS Prioritization on UPF	2023.04	The Wireless Priority Services (WPS) feature provides finer control for priority handling over multiple interfaces. UPF supports WPS services based on the message priority indicated by SMF. The configured priority value set on SMF will be sent to UPF over N4 as part of the PFCP header.

- Feature Description, on page 603
- How it Works, on page 604
- Message Prioritization for Overload Scenarios, on page 607
- Configuring WPS Prioritization, on page 607
- OAM Support, on page 609

Feature Description

The Wireless Priority Services (WPS) feature provides finer control for priority handling over multiple interfaces. UPF supports WPS based on the message priority indicated by SMF.

The priority value set on SMF will be sent to UPF over N4 as part of the PFCP header. The message priority values from SMF and cnSGWc are currently supported for Emergency, VoLTE, and non-VoLTE IMS sessions.

Note WPS supports only N4 (SMF) and Sxa (cnSGW) terminations on UPF. WPS does not support Sxb (PGW-C) and Sxa (SGW-C) terminations in this release.

For information on WPS prioritization support on RCM, see the *UCC 5G RCM Configuration and Administration Guide*.

For information on WPS prioritization support on SMF, see the *UCC 5G SMF Configuration and Administration Guide*.

How it Works

This section describes how WPS prioritization works in UPF.

Priority Handling

The session-priority-profile maps the session-priority values from 0 to 15 to precedence values from 1 to 4. It is used internally in modules like RCM and data-queue distribution.

SR and ICSR

When the feature is not enabled, any session with a valid session-priority received over N4/SX will be prioritized for SR/ICSR as per earlier design.

When the feature enabled, there is a change such that only the Emergency and WPS sessions will be prioritized for SR/ICSR. During UPF downgrade, the priority values are recovered with an offset of +1.

Data Queue

The following table describes the traffic distribution over data queues based on session precedence:

Queue #	Details
3 (Highest priority)	Reserved for Sx
2	Used for Precedence 1,2 sessions
1	Used for Precedence 3,4 sessions
0 (Lowest priority)	Used for Data/Non-prioritized sessions

The default precedence values will be used when:

- Only session-type is configured and precedence is not defined for session-priority
- No priority value is configured in the session-priority profile
- WPS feature is enabled with the **default** session-priority profile

Note

- If the feature is disabled, UPF uses N-1 support prioritization.
- For session recovery of combo calls, the highest priority value received on Sxa or N4 will be used for both sessions.

Mapping to Session Type

UPF maps the session priority values (WPS, Emergency, and IMS) received over N4 to the session types using the "associate session-priority-profile" CLI configuration.

For VoLTE non-active sessions, mapping is not available. For VoLTE active sessions, the priority value 'ims' will be used, while the next priority value (+1) is implicitly used for VoLTE non-active sessions. Hence, "ims+1" priority should not be configured in the session-priority profile.

Version Compatibility

1. For the latest version (X) of RCM and a prior version (X-1) of UPF:
   a. RCM will support the previous UPF version (X-N), as the priority set is extended without affecting the previously supported set of priorities (1, 2, 3).
   b. During switchover, the flush sequence will be completely based on the priority received from the UPF, which was not the case in earlier versions.
2. For the prior version (X-1) of RCM and latest version (X) of UPF:
   a. RCM will not support the extended set of priorities, and the checkpoints received for priority 4 will be dropped.
   b. During switchover, calls for priority 1, 2, and 3 can be restored.

Control-Plane Compatibility

To enable WPS prioritization, upgrade the UPF to the new configuration. Prioritization will not be guaranteed in the following scenarios when:

- SMF/CP is not upgraded
- Priority values are sent based on N-1 support
- WPS prioritization is enabled on UPF

Based on the priority values received from the Control Plane, the following combinations will work:

S. No.	SMF / cnSGW / PGW-C / SGW-C Version	UPF Version	WPS prioritization enabled on UPF	UPF Behavior
1	N-1 Or N without config	N-1	N/A	No WPS prioritization - same as previous behavior.
2	N-1 Or N without config	N	No	No WPS prioritization - same as previous behavior.

S. No.	SMF / cnSGW / PGW-C / SGW-C Version	UPF Version	WPS prioritization enabled on UPF	UPF Behavior
3	N-1 Or N without config	N	Yes	Session prioritization will not be guaranteed.
4	N (Config present)	N-1	N/A	This is not recommended as session prioritization will not be guaranteed. RCM checkpoints will be dropped for priority>3.
5	N (Config present)	N	Yes	WPS prioritization works.
6	N (Config present)	N	No	This is not recommended as session prioritization will not be guaranteed.

Limitations

The WPS Prioritization feature has the following limitations:

- In a multi-Sx/N4 scenario, all Control Plane nodes must use the same process to send MP values. This applies to SMF/cnSGW and PGW-C/SGW-C working with the same UPF.
- The sessions recovered during the UPF software upgrade will retain their prioritization values. As a result, after the UPF upgrade between N-1 and N, the prioritization for existing sessions will not work as expected, and the show CLI output will be undetermined.
- In the case of a configuration change at the SMF that changes existing priority values or a configuration change at the UPF that changes the session type, the session type and priority are not re-evaluated. As a result of such a configuration change, incorrect statistics updates may occur.
- Due to any mid-session priority changes, some of the statistics may be inconsistent.

Message Prioritization for Overload Scenarios

Table 133: Feature History

Feature Name	Release Information	Description
Handling PFCP Messages Using Message Prioritization during an Overload Scenario	2024.04.1	This feature allows the UPF to use message prioritization to ensure uninterrupted receival of incoming PFCP messages during an overload scenario or self-protection mode. Message Prioritization allows the network operator to define the message priority using the configuration to throttle the PFCP messages in overload scenarios for a WPS session. **Command Enhanced**: **session-priority-profile** *spp_name* **priority** *priority_value* **type { wps | emergency | ims } { throttle | precedence** *precedence_value* **}** **Default Setting:** Disabled-Configuration required to enable.

Message Priority for PFCP Messages

During an overload condition, UPF handles the incoming PFCP messages using Message Priority. UPF receives the message priority bits in the PFCP signaling, based on the operator policy. UPF sets the same message priority in the outgoing response message that it receives in the incoming request message.

UPF compares this message priority with the message priority defined in the configuration. UPF then throttles or exempts from throttling according to the configuration in the overload or self-protection mode.

UPF exempts the priority traffic from throttling unless it is impossible to achieve the requested traffic reduction without throttling the priority traffic as well. By default, UPF exempts the priority PFCP messages from throttling. However, you can throttle the priority PFCP sessions by defining specifically in the configuration.

To throttle the priority PFCP messages, see the Mapping the Session Priority Profile section.

Configuring WPS Prioritization

This section describes how to configure the WPS Prioritization feature.

Enabling WPS Prioritization

Use the following configuration to enable or disable WPS prioritization support and associate a session-priority-profile with the user-plane-service:

```
config
  context context_name
    user-plane-service up_service_name
      [ no ] associate session-priority-profile { name | default } [
name_string ]
      end
```

NOTES:

- **associate session-priority-profile { name | default } [** *name_string* **]**: Enables WPS prioritization support.
  - **name**: Specify the name of the session priority profile to be associated with the User Plane service.
  - **default**: Specify the default precedence value associated with the session priority.
  - *name_string* : Specify the name string associated with the session priority.
- **no associate session-priority-profile** : Disables WPS prioritization support.
- On a given UPF, you can configure a total number of 16 session-priority profiles. Only one session-priority profile can be associated with a user-plane service.
- When a session does not receive any message-priority within the N4/Sx messages, it is classified as a 'Normal' session with a 'precedence' value set to 0.
- If a priority is not specified in the session-priority-profile configuration, the session-type will be shown as 'Unclassified' and uses the default1 precedence value.
- When there is no session-priority-profile association in the user-plane-service, UPF utilizes the existing (N-1) behavior with fixed priority values, including Emergency, IMS-VoLTE, and IMS-nonVoLTE.

Mapping the Session Priority Profile

Use the following configuration to configure the session priority value with a session type and precedence value.

```
config
   session-priority-profile spp_name
     priority priority_value type { wps | emergency | ims } [ throttle |
precedence precedence_value ]
     end
```

NOTES:

- **session-priority-profile**: Specify the session priority profile name.
- **priority**: Specify the priority value as an integer in the range of 0–15.
- **type { wps | emergency | ims }**: Specify the session type.
- **throttle**: Defines the behavior for flows of priority sessions during overload/self-protection state to be throttled. The default value of this keyword is Enabled, if the priority is not configured.
- **precedence**: Specify the precedence value as an integer in the range of 1–4.
- The session-priority profile must not be configured with the priority value next to the one assigned for **ims**.

OAM Support

Show Commands Support

This section provides information about show commands and their outputs in support of this feature.

show subscribers user-plane-only full all

The output of this CLI command is enhanced to display the session-type value.

The following is a sample output of this command:

```
show subscribers user-plane-only full all | more

  Local SEID      : [0x0004000000000000] 1125899906842624
  Remote SEID     : [0x0004000000000002] 1125899906842626
  State           : Connected
  Connect Time    : Thu Jun  8 09:36:49 2023
Idle time       : 00h00m11s
  Access Type: uplane-ipv4              Network Type: IP
  user-plane-service-name: user-plane-service
  active-service-scheme-name:
  Callid: 00004e21
  Rulebase: prepaid
  Interface Type: Sxb
  eMPS Session: Yes
  eMPS Session Priority: 6
  Session-Type: wps
  Precedence-Order: 1
  Data Queue: 2
```

show user-plane-service all

The output of this CLI command is enhanced to display the Session Priority Profile name.

The following is a sample output of this command:

```
show user-plane-service all

Service name : user-plane-service
Service-Id : 7
Context : ingress
Status : STARTED
PGW Ingress GTPU Service : sx-gtpu-service
…
Nf-Instance-Id : Not defined
Session Priority Profile          : sp1
```

show user-plane-service session-priority-profile name | all

The output of this CLI command displays the session priority profile name, along with type and precedence value per priority.

The following is a sample output of this command:

```
show user-plane-service session-priority-profile name spp1

  Session Priority Profile Name: spp1
```

```
  Priority     Type            Precedence
  -------------------------------------
  0            unclassified    1
  1            wps             1
  2            unclassified    1
  3            emergency       2
  4            unclassified    2
  5            ims             3
  6            ims             3
  7            unclassified    2
  8            unclassified    3
  9            unclassified    3
  10           unclassified    3
  11           unclassified    3
  12           unclassified    4
  13           unclassified    4
  14           unclassified    4
  15           unclassified    4

Total session priority profile(s) found: 1
```

show user-plane-service statistics all

The output of this command is enhanced to display the number of WPS sessions.

The following is a sample output of this command:

```
show user-plane-service statistics all
…
  eMPS PDNs Total:
    Active:                          0        Setup:                          0
    Released:                        0        Rejected:                       0

  WPS PDNs Total:
    Active:                          0        Setup:                          0
    Released:                        0        Rejected:                       0
```

show user-plane-service statistics session-priority-profile

The output of this CLI command displays the number of prioritized and non-prioritized sessions. This command also displays prioritized sessions based on the ‘type’ configuration in the ‘session-priority-profile’.

The following is a sample output of this command:

```
show user-plane-service statistics session-priority-profile

  Total Non-Prioritized (Normal) sessions: 0

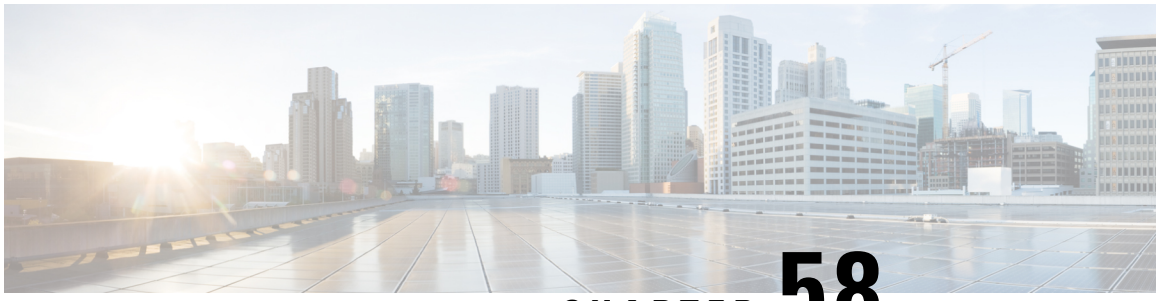
  Total Prioritized sessions: 121
    WPS sessions            Total: 11          Active: 11
      Precedence            Config match       Default match
      ----------           --------------     ---------------
        1                              1                   0
        2                              0                  10
        3                              0                   0
        4                              0                   0

    Emergency sessions      Total: 110         Active: 110
      Precedence            Config match       Default match
      ----------           --------------     ---------------
        1                              0                  10
        2                              0                   0
        3                              0                 100
```

```
    4                                0                    0

IMS sessions            Total: 0             Active: 0
  Precedence            Config match         Default match
  ----------           --------------       ---------------
    1                                0                    0
    2                                0                    0
    3                                0                    0
    4                                0                    0

Unclassified sessions   Total: 0             Active: 0
  Precedence            Config match         Default match
  ----------           --------------       ---------------
    1                                0                    0
    2                                0                    0
    3                                0                    0
    4                                0                    0
```



CHAPTER 58

X-Header Insertion and Encryption

- Feature Summary and Revision History, on page 613
- Feature Description, on page 614
- How it Works, on page 614
- Configuring X-Header Insertion and Encryption, on page 615
- Monitoring and Troubleshooting X-Header Insertion and Encryption Configuration, on page 617
- Anti-Spoofing Support, on page 618

Feature Summary and Revision History

Summary Data

Table 134: Summary Data

Applicable Product(s) or Functional Area	5G-UPF
Applicable Platform(s)	VPC-SI
Feature Default Setting	Disabled – Configuration Required
Related Changes in This Release	Not Applicable
Related Documentation	*UCC 5G UPF Configuration and Administration Guide*

Revision History

Revision Details	Release
Added support for the anti-spoofing functionality through X-header enrichment.	2023.02.0
First introduced.	2020.02.x

Feature Description

The process of X-header insertion and encryption is collectively known as Header Enrichment (HE). The UPF supports HE to allow deletion and modification of the existing X-header fields in the HTTP header of the URL requests.

This feature enables appending the headers to HTTP or WSP GET and POST Request packets, and HTTP Response packets for use by end applications, such as mobile advertisement insertion (MSISDN, IMSI, IP address, user-customizable, and so on).

How it Works

HE includes two operations:

1. Header Insertion
2. Header Encryption

X-Header Insertion

Extension header (X-header) fields are fields that are added to protocol headers for specific purposes such as user identification, device detection, etc. The X-header mechanism allows additional entity-header fields to be defined without changing the protocol. However, these fields are assumed to be unrecognizable by the recipients.

X-header insertion inserts X-headers in HTTP or WSP GET and POST Request packets and HTTP Response packets. You can configure rules to insert X-headers in HTTP or WSP Request and HTTP Response packets. The charging-action associated with the rules contain the list of X-headers to be inserted in the packets.

Note

- Flows for which the X-header is inserted in a packet are not offloaded.
- With the X-header configuration, all TCP OOO packets irrespective of transmit order CLI, will be buffered and sent after re-ordering.

X-Header Encryption

X-header encryption enables encryption of X-header fields before insertion. If X-header insertion has already happened for an IP flow (because of any X-header format) and the current charging-action has the first-request-only flag set, the X-header insertion will not happen for that format. If the first-request-only flag is not set in a charging-action, the insertion for that X-header format continues happening in other suitable packets of that IP flow.

Changes to the X-header format configuration do not trigger re-encryption for existing calls. The changed configuration will however, be applicable for new calls. The changed configuration will also apply at the time of the next re-encryption, to the existing calls for which the re-encryption timeout is specified. If encryption

is enabled for a parameter while data is flowing, since its encrypted value will not be available, insertion of that parameter stops.

Note This feature does not support flow recovery.

Limitations

This section highlights the limitations existing in the current release:

- X-header insertion in the Response packet is not supported.
- X-header encryption with RSA and RC4MD5 is supported, but not supported with AES.
- Monitor protocol for X-header is not supported.
- Insertion of the following X-header fields is not supported in a packet:
  - QoS
  - UIDH
  - Customer ID
  - Hash Value
  - Time of the Day
  - RADIUS String
  - Session-Id
  - Congestion Level
  - User-Profile
- Parsing and buffering of the HTTP header will be done only if the rulebase has any x-header format configured with x-header fields explicitly set to delete-existing. For operators not using x-header delete-existing, no overhead for parsing or buffering will be seen.

Configuring X-Header Insertion and Encryption

Configuring X-Header Insertion

Step 1	Create or configure a ruledef to identify the HTTP packets in which the X-headers must be inserted.
Step 2	Create or configure a rulebase and configure the charging-action that inserts the X-header fields into the HTTP packets.
Step 3	Create the X-header format as described in *Creating the X-Header Format*.

Step 4	Configure the X-header format based on the message type in the charging action, as described in *Configuring the X-Header Format*.

Creating the X-Header Format

To create an X-header format, use the following configuration:

```
configure
   active-charging service ecs_service_name
      xheader-format xheader_format_name
      end
```

Configuring the X-Header Format

To configure an X-header format, use the following configuration:

```
configure
   active-charging service ecs_service_name
      xheader-format xheader_format_name
        insert xheader_field_name string-constant xheader_field_value | variable
{ bearer { 3gpp { apn | charging-characteristics | charging-id | imei |
imsi | qos | rat-type | s-mcc-mnc | sgsn-address } | acr | customer-id |
 ggsn-address | mdn | msisdn-no-cc | radius-string |
radius-calling-station-id | session-id | sn-rulebase |
subscriber-ip-address | username } [ encrypt ] [ delete-existing ] | http
 { host | url } }
        end
```

Configuring X-Header Encryption

Step 1	Configure the X-header insertion as described in *Configuring X-Header Insertion*.
Step 2	Create or configure a rulebase, and the encryption certificate to use and the re-encryption parameter as described in *Configuring X-Header Encryption*.
Step 3	Configure the encryption certificate to use as described in *Configuring Encryption Certificate*.

Configuring X-Header Encryption

To configure X-header encryption, use the following configuration:

```
configure
   active-charging service ecs_service_name
      rulebase rulebase_name
         xheader-encryption certificate-name certificate_name
         xheader-encryption re-encryption period re-encryption_period
         end
```

NOTES:

- This configuration enables X-header encryption for all subscribers using the specified rulebase.

- If the certificate is removed, ECS continues using the copy that it has. The copy is set free once the certificate name is removed from the rulebase.
- Changes to x-header format configuration won't trigger re-encryption for existing calls. The changed configuration will however, be applicable for new calls. The changed configuration will also apply at the next reencryption time to those existing calls for which reencryption timeout is specified. If encryption is enabled for a parameter while data is flowing, since its encrypted value won't be available, insertion of that parameter stops.

Configuring Encryption Certificate

To configure the encryption certificate, use the following configuration:

```
configure
  certificate name certificate_name pem { { data pem_certificate_data private-key
 pem [ encrypted ] data pem_pvt_key } | { url url private-key pem { [
encrypted ] data pem_pvt_key | url url } }
  end
```

Verifying X-header Insertion and Encryption Configuration

Use the following command in the Exec Mode to verify your configuration:

xheader-format *xheader_format_name*

Monitoring and Troubleshooting X-Header Insertion and Encryption Configuration

show active-charging charging-action statistics name

The output of this command displays statistics for X-header information.

- X-header information:
  - XHeader Bytes Injected
  - XHeader Pkts Injected
  - IP Frags consumed by XHeader
  - XHeader Bytes Removed
  - XHeader Pkts Removed

show active-charging rulebase statistics name

The output of this command displays the header enrichment statistics.

- HTTP header buffering limit reached

Anti-Spoofing Support

Feature Description

The UPF supports spoofing detection and mitigation using header enrichment. The anti-spoofing functionality allows the operators and their subscribers from spoofing threats posed by malicious UE devices. This feature detects the fraudulent activities when an external portal is used for subscriber or content authorization.

The anti-spoofing functionality is supported only for HTTP Requests. This functionality allows deletion and modification of the existing X-header fields in the HTTP header of the URL request.

If the user-generated header name is similar to the header name used in x-header format, you can search, delete, and modify the X-header fields. The X-header field names are case-insensitive. For example, if x-MSISDN and x-msisdn fields are present in an incoming HTTP request, both fields are processed.

If the user HTTP header already has the fields to be inserted, the corresponding values for those fields are replaced with the modified values from the gateway at the end of the header. In case of multiple entries for one field, which is already present in the header, a single instance of the field is selected along with the value inserted from the gateway.

In case of multiple HTTP GET requests, the packet header replaced with X-header fields for all GET requests. If encryption is enabled, it is performed after inserting the X-header.

Limitations

The following are the known limitations of the anti-spoofing feature:

- If an HTTP GET or POST header is not complete in 3 packets, the anti-spoofing is performed only for the last packet in which the header completes. Otherwise, the packets are sent out without HE.
- Though fields without having 'x-' in the field name are allowed, we do not delete the extension header fields that do not start with x-.

Enabling Anti-Spoofing in X-Header

To enable anti-spoofing in the X-header, use the following configuration:

```
configure
  active-charging service acs_service_name
    xheader-format xheader_format_name
      insert xheader_field_name variable { bearer subscriber-ip-address } [
 encrypt ] [ delete-existing ]
      exit
```

NOTES:

- **delete-existing**: This command enables detection of spoofing in an X-header file.
- This feature is disabled by default.
- A maximum number of 10 X-header fields can be configured with the **delete-existing** CLI keyword, which also includes fields configured only for X-header insertion.

OAM Support

This section provides information regarding show commands and their outputs for the anti-spoofing feature.

show user-plane-service statistics charging-action [name | all]

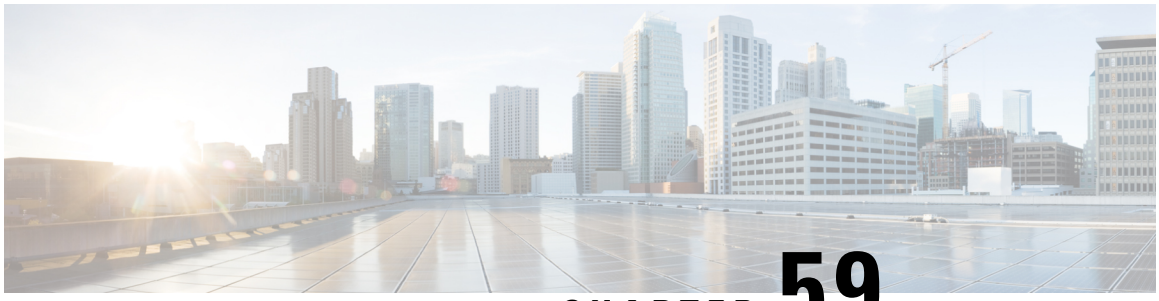
The output of this command displays the user plane statistics for charging action.

- X-header information:
  - XHeader Bytes Injected
  - XHeader Pkts Injected
  - XHeader Bytes Removed
  - XHeader Pkts Removed

show user-plane-service statistics rulebase [name | all]

The output of this command displays user plane statistics for rulebase.

- Header Enrichment stats:
  - HTTP header buffering limit reached



CHAPTER 59

UPF Troubleshooting Information

This chapter covers the following topics related to monitoring and troubleshooting the UPF features:

- Debug Logging , on page 621
- Monitoring CLI, on page 622
- Monitoring Protocol, on page 623
- RAT Type-based Statistics, on page 623
- Subscriber Level CLI, on page 627
- VPP Statistics, on page 635
- SNMP Support, on page 639
- Troubleshooting UPF Features, on page 640
- Configure Stats Profile for QCI Drop Statistics, on page 642
- NSO configuration considerations, on page 643
- Per-QCI Drop statistics commands, on page 643

Debug Logging

Important The debug logging CLIs must be enabled with the help of System Administrator. Enabling debug logging CLIs can be resource intensive.

Use the following debug CLIs as required:

- **logg filter active facility sx level debug**
- **logg filter active facility user-data level debug**
- **logg filter active facility sessmgr level debug**
- **logg filter active facility uplane level debug**
- **logg filter active facility egtpc level debug**
- **logg filter active facility gtpu level debug**
- **logg filter active facility egtpu level debug**
- **logg filter active facility gtpumgr level debug**

- **logg filter active facility sxdemux level debug**
- **logg filter active facility user-l3tunnel level debug**
- **logg filter active facility aaamgr level debug**
- **logg filter active facility vpp level debug**
- **logg filter active facility dpath level debug**
- **logg active pdu-verbosity 5**
- **logg syslog** *ip_address* **facility** *facilities* **event-verbosity { min | concise | full }**

Monitoring CLI

Subscriber Level Message

Use the **mon sub callid** CLI command for subscriber level message.

Resource Tracking

Use the **show task resources facility** *sessmgr* **all** CLI command to track the CPU/Memory for PROCLET.

Service Status

Use the **show service all** CLI command to check the service status.

Sx Peer Status

Use the **show sx peers** CLI command to check the Sx peer status.

This output of this CLI is enhanced to display the auto config success, config push start time, and config push end time for a Sx peer.

```
[local]laas-si-setup# show sx peers wide
         ID                                                          Restart

|||||||  |                                        Recovery          |   Current   Max
  Peer   Config  Auto-Config  Config Push  Config Push
vvvvvvv  v    Group Name  Node ID      Peer ID     Timestamp         v  Sessions  Session
 State  Failures  Success    Start Time    End Time
------  ----  ----------- ---------   ---------  ------------------ ---- ------- -------
------  -------  --------    ----------   -----------
CAAXDXN  9    CPGROUP11   20.20.20.81  33554433 2025-05-14:05:36:46  0     0        0
 NONE     0         0           NA           NA
CAAXDXN  22   CPGROUP21   20.20.20.101 50331649 2025-05-14:05:36:45  0     0        0
 NONE     0         0           NA           NA

Total Peers:    2
```

GTP-U Peer Status

Use the **show gtpu peers** CLI command to check the GTP-U peer status.

Monitoring Protocol

When using the monitor protocol command, enable option 49 for PFCP, and option 26 for GTP-U.

RAT Type-based Statistics

The RAT Type-based Statistics feature equip users to view data statistics segregated by RAT Type in UPF.

RAT Type-based data statistics in UPF maintains separate buckets. These buckets are created at Session Manager instance level. Bucket is assigned to a subscriber at the time of call-setup, based on RAT Type IE received in “Subscriber-Parameters”. If the IE is not received, “Unknown” RAT Type bucket is assigned to that subscriber. During the session, if UPF receives a new RAT Type for a subscriber, the bucket is changed accordingly.



Important Data statistics are not checkpointed and lost during Session Recovery/ICSR. Only “Current-Subscriber” statistics are recalculated after recovery (during the time of call-audit).

Show Command and Output

The following CLI command displays node-level RAT statistics for UPF: **show user-plane-service statistics rat { 5g-nr | all | eutran | unknown | wlan }**

NOTES:

- **5g-nr**: Displays the data statistics for 5G NR subscribers.
- **all**: Displays the data statistics for all RAT Type subscribers.
- **eutran**: Displays the data statistics for EUTRAN subscribers.
- **unknown**: Displays the data statistics for subscribers of unknown RAT type.
- **wlan**: Displays the data statistics for WLAN subscribers.

Statistics

The following table provides description of each field.

Table 135: show user-plane-service statistics rat all

Field	Description
Current Subscribers	
5G NR	Specifies the total number of current 5G NR subscribers.
EUTRAN	Specifies the total number of current EUTRAN subscribers.
WLAN	Specifies the total number of current WLAN subscribers.

Field	Description
Unknown	Specifies the total number of current subscribers of unknown RAT type.
Data Statistics	
5G NR	Specifies the data statistics for 5G NR subscribers.
Uplink	Specifies data statistics for 5G NR subscribers in uplink direction.
Total Pkts	Specifies the total number of uplink packets for 5G NR subscribers.
Total Bytes	Specifies the total number of uplink bytes for 5G NR subscribers.
Total Dropped Pkts	Specifies the total number of uplink packets dropped for 5G NR subscribers.
Total Dropped Bytes	Specifies the total number of uplink bytes dropped for 5G NR subscribers.
Downlink	Specifies data statistics for 5G NR subscribers in downlink direction.
Total Pkts	Specifies the total number of downlink packets for 5G NR subscribers.
Total Bytes	Specifies the total number of downlink bytes for 5G NR subscribers.
Total Dropped Pkts	Specifies the total number of downlink packets dropped for 5G NR subscribers.
Total Dropped Bytes	Specifies the total number of downlink bytes dropped for 5G NR subscribers.
EUTRAN	Specifies the data statistics for EUTRAN subscribers.
Uplink	Specifies data statistics for EUTRAN subscribers in uplink direction.
Total Pkts	Specifies the total number of uplink packets for EUTRAN subscribers.
Total Bytes	Specifies the total number of uplink bytes for EUTRAN subscribers.
Total Dropped Pkts	Specifies the total number of uplink packets dropped for EUTRAN subscribers.
Total Dropped Bytes	Specifies the total number of uplink bytes dropped for EUTRAN subscribers.
Downlink	Specifies data statistics for EUTRAN subscribers in downlink direction.
Total Pkts	Specifies the total number of downlink packets for EUTRAN subscribers.
Total Bytes	Specifies the total number of downlink bytes for EUTRAN subscribers.
Total Dropped Pkts	Specifies the total number of downlink packets dropped for EUTRAN subscribers.
Total Dropped Bytes	Specifies the total number of downlink bytes dropped for EUTRAN subscribers.
WLAN	Specifies the data statistics for WLAN subscribers.

Field	Description
Uplink	Specifies data statistics for WLAN subscribers in uplink direction.
Total Pkts	Specifies the total number of uplink packets for WLAN subscribers.
Total Bytes	Specifies the total number of uplink bytes for WLAN subscribers.
Total Dropped Pkts	Specifies the total number of uplink packets dropped for WLAN subscribers.
Total Dropped Bytes	Specifies the total number of uplink bytes dropped for WLAN subscribers.
Downlink	Specifies data statistics for WLAN subscribers in downlink direction.
Total Pkts	Specifies the total number of downlink packets for WLAN subscribers.
Total Bytes	Specifies the total number of downlink bytes for WLAN subscribers.
Total Dropped Pkts	Specifies the total number of downlink packets dropped for WLAN subscribers.
Total Dropped Bytes	Specifies the total number of downlink bytes dropped for WLAN subscribers.
Unknown	Specifies the data statistics for subscribers of unknown RAT type.
Uplink	Specifies data statistics for unknown RAT type subscribers in uplink direction.
Total Pkts	Specifies the total number of uplink packets for unknown RAT type subscribers.
Total Bytes	Specifies the total number of uplink bytes for unknown RAT type subscribers.
Total Dropped Pkts	Specifies the total number of uplink packets dropped for unknown RAT type subscribers.
Total Dropped Bytes	Specifies the total number of uplink bytes dropped for unknown RAT type subscribers.
Downlink	Specifies data statistics for unknown RAT type subscribers in downlink direction.
Total Pkts	Specifies the total number of downlink packets for unknown RAT type subscribers.
Total Bytes	Specifies the total number of downlink bytes for unknown RAT type subscribers.
Total Dropped Pkts	Specifies the total number of downlink packets dropped for unknown RAT type subscribers.
Total Dropped Bytes	Specifies the total number of downlink bytes dropped for unknown RAT type subscribers.

Bulk Statistics

The following bulk statistics are included in the User Plane Service schema to track RAT Type-based data statistics events.

Table 136: show bulkstats variables user-plane-service

Variable Name	Data Type	Key	Counter Type
vpnname	String	1	Info
vpnid	Int32	1	Info
servname	String	1	Info
servid	Int32	1	Info
curr-pdn-rat-eutran	Int64	0	Gauge
curr-pdn-rat-5g-nr	Int64	0	Gauge
curr-pdn-rat-wlan	Int64	0	Gauge
curr-pdn-rat-unknown	Int64	0	Gauge
uplink-total-pkts-pdn-rat-eutran	Int64	0	Counter
uplink-total-bytes-pdn-rat-eutran	Int64	0	Counter
uplink-total-pkts-dropped-pdn-rat-eutran	Int64	0	Counter
uplink-total-bytes-dropped-pdn-rat-eutran	Int64	0	Counter
downlink-total-pkts-pdn-rat-eutran	Int64	0	Counter
downlink-total-bytes-pdn-rat-eutran	Int64	0	Counter
downlink-total-pkts-dropped-pdn-rat-eutran	Int64	0	Counter
downlink-total-bytes-dropped-pdn-rat-eutran	Int64	0	Counter
uplink-total-pkts-pdn-rat-5g-nr	Int64	0	Counter
uplink-total-bytes-pdn-rat-5g-nr	Int64	0	Counter
uplink-total-pkts-dropped-pdn-rat-5g-nr	Int64	0	Counter
uplink-total-bytes-dropped-pdn-rat-5g-nr	Int64	0	Counter
downlink-total-pkts-pdn-rat-5g-nr	Int64	0	Counter
downlink-total-bytes-pdn-rat-5g-nr	Int64	0	Counter
downlink-total-pkts-dropped-pdn-rat-5g-nr	Int64	0	Counter
downlink-total-bytes-dropped-pdn-rat-5g-nr	Int64	0	Counter

Variable Name	Data Type	Key	Counter Type
uplink-total-pkts-pdn-rat-wlan	Int64	0	Counter
uplink-total-bytes-pdn-rat-wlan	Int64	0	Counter
uplink-total-pkts-dropped-pdn-rat-wlan	Int64	0	Counter
uplink-total-bytes-dropped-pdn-rat-wlan	Int64	0	Counter
downlink-total-pkts-pdn-rat-wlan	Int64	0	Counter
downlink-total-bytes-pdn-rat-wlan	Int64	0	Counter
downlink-total-pkts-dropped-pdn-rat-wlan	Int64	0	Counter
downlink-total-bytes-dropped-pdn-rat-wlan	Int64	0	Counter
uplink-total-pkts-pdn-rat-unknown	Int64	0	Counter
uplink-total-bytes-pdn-rat-unknown	Int64	0	Counter
uplink-total-pkts-dropped-pdn-rat-unknown	Int64	0	Counter
uplink-total-bytes-dropped-pdn-rat-unknown	Int64	0	Counter
downlink-total-pkts-pdn-rat-unknown	Int64	0	Counter
downlink-total-bytes-pdn-rat-unknown	Int64	0	Counter
downlink-total-pkts-dropped-pdn-rat-unknown	Int64	0	Counter
downlink-total-bytes-dropped-pdn-rat-unknown	Int64	0	Counter

Subscriber Level CLI

Use the following subscriber level CLIs as required:

- **show subscribers user-plane-only full all**
- **show subscribers user-plane-only full msid** *msid_value*
- **show subscribers user-plane-only { seid** *seid_value* **| callid** *callid_value* **} pdr all**
- **show subscribers user-plane-only { seid** *seid_value* **| callid** *callid_value* **} far all**
- **show subscribers user-plane-only { seid** *seid_value* **| callid** *callid_value* **} qer all**
- **show subscribers user-plane-only { seid** *seid_value* **| callid** *callid_value* **} urr all**
- **show subscribers user-plane-only { seid** *seid_value* **| callid** *callid_value* **} bar all**
- **show subscribers user-plane-only { seid** *seid_value* **| callid** *callid_value* **} pdr full all**
- **show subscribers user-plane-only { seid** *seid_value* **| callid** *callid_value* **} urr full all**

- **show subscribers user-plane-only { seid** *seid_value* **| callid** *callid_value* **} far full all**
- **show subscribers user-plane-only { seid** *seid_value* **| callid** *callid_value* **} qer full all**
- **show subscribers user-plane-only { seid** *seid_value* **| callid** *callid_value* **} bar full all**
- **show subscribers user-plane-only { seid** *seid_value* **| callid** *callid_value* **} pdr id** *id_value*
- **show subscribers user-plane-only { seid** *seid_value* **| callid** *callid_value* **} flows full**
- **show subscribers user-plane-only { seid** *seid_value* **| callid** *callid_value* **} bli full all**

show subscribers user-plane-only full msid *msid_value*

The **show subscribers user-plane-only full msid** *msid_value* command displays the specific user information at instance-level and session-level for data collection during troubleshooting.

The following is a sample output of this command:

```
show subscribers user-plane-only full msid 404005123456789

  Local SEID      : [0x0004000000000000] 1125899906842624
  Remote SEID     : [0x0004000000000007] 1125899906842631
State           : Connected
  Connect Time    : Thu Sep 15 20:40:09 2022
  Idle time       : 00h11m25s
  Access Type: uplane-ipv4               Network Type: IP
  user-plane-service-name: UP21
  active-service-scheme-name:
  Callid: 00004e21
  Rulebase: cisco
  Interface Type: Sxb
  eMPS Session: No
  eMPS Session Priority: 0
  Card/Cpu: 1/0                          Sessmgr Instance: 1
  IP address: 11.11.12.0
  Next Hop Ip Address:
  Source context: EPC2-UP                Destination context: ISP1-UP
  PDN-Instance: starent.com
  User-plane-Sx-addr: 20.20.20.106
  Control-plane-Sx-addr: 20.20.20.101
  Number of associated PDRs : 0
  Number of associated ADC PDRs : 0
  Number of associated FARs : 0
  Number of associated QERs : 1
  Number of associated BARs : 0
  Number of associated URRs : 0
  Uplink APN AMBR (bps) : 200000       Downlink APN AMBR (bps) : 200000
  active input acl: IPV4ACL            active output acl: IPV4ACL
  active input ipv6 acl: n/a             active output ipv6 acl: n/a
  Bandwidth Policy: bw_policy1
  FW-and-NAT Policy: n/a
  FW-and-NAT Policy ID: n/a
  Firewall Policy IPv4: n/a
  Firewall Policy IPv6: n/a
  NAT Policy NAT44: n/a
  NAT Policy NAT64: n/a
  Converged Session: No                 Converged Peer Callid:      n/a
  Visited Call: No
  Subscriber Parameters:
  IMSI: 404005123456789
  IMEI: 112233445566778
```

```
  MSISDN: 9890098900
  Charging Characteristics: 0x3412
  Rat Type: 6
  MCC MNC: 123765
  SGSN Address: 20.20.20.61
  ULI: TAI=12546312ab00
  Congestion Level: 0
  GGSN Address: 20.20.20.101
  User-Name: 9890098900@starent.com
  Session-ID: 1414146500000008
  input pkts: 0                                  output pkts: 0
  input bytes: 0                                 output bytes: 0
  input bytes dropped: 0                         output bytes dropped: 0
  input pkts dropped: 0                          output pkts dropped: 0
  CF Buffered Uplink Packets: 0                  CF Buffered Downlink Packets: 0

  CF Buffered Uplink Bytes: 0                    CF Buffered Downlink Bytes: 0
  Uplink Packets in Buffer: 0                    Downlink Packets in Buffer: 0
  Buff Over-limit Uplink Pkts: 0                 Buff Over-limit Downlink Pkts: 0

  DDN buffered pkts : 0                          DDN buffered bytes : 0
  DDN buffer overflow drop pkts : 0              DDN buffer overflow drop bytes : 0

  pk rate from user(bps): 0                      pk rate to user(bps): 0
  ave rate from user(bps): 0                     ave rate to user(bps): 0
  sust rate from user(bps): 0                    sust rate to user(bps): 0
  pk rate from user(pps): 0                      pk rate to user(pps): 0
  ave rate from user(pps): 0                     ave rate to user(pps): 0
  sust rate from user(pps): 0                    sust rate to user(pps): 0
  ipv4 bad hdr: 0                                ipv4 ttl exceeded: 0
ipv4 fragments sent: 0                         ipv4 could not fragment: 0
  ipv4 bad length trim: 0
  input pkts dropped (0 mbr): 0                  output pkts dropped (0 mbr): 0
  ipv4 input acl drop: 0                         ipv4 output acl drop: 0
  ipv6 input acl drop: 0                         ipv6 output acl drop: 0
  ip source violations: 0                        ipv4 output no-flow drop: 0
  ipv6 bad hdr: 0                                ipv6 bad length trim: 0
  ipv4 icmp packets dropped: 0
  APN AMBR Uplink Pkts Drop: 0                    APN AMBR Downlink Pkts Drop: 0
  APN AMBR Uplink Bytes Drop: 0                   APN AMBR Downlink Bytes Drop: 0

  APN AMBR Uplink Pkts IP pref lowered: 0         APN AMBR Downlink Pkts IP pref lowered:
 0
  APN AMBR Uplink Bytes IP pref lowered: 0        APN AMBR Downlink Bytes IP pref lowered:
 0
  ITC Uplink Pkts Drop: 0                         ITC Downlink Pkts Drop: 0
  ITC Uplink Bytes Drop: 0                        ITC Downlink Bytes Drop: 0
  ITC Uplink Pkts IP pref lowered: 0              ITC Downlink Pkts IP pref lowered: 0

  ITC Uplink Bytes IP pref lowered: 0             ITC Downlink Bytes IP pref lowered: 0

  ITC Terminated Flows:: 0                        ITC Redirected Flows: 0
  Flow Action Terminated Flows: 0                 Flow Action Redirected Flows: 0

  ToS marked Uplink Pkts: 0                       ToS marked Downlink Pkts: 0
  CC Dropped Uplink Pkts: 0                       CC Dropped Downlink Pkts: 0
  CC Dropped Uplink bytes: 0                      CC Dropped Downlink Bytes: 0
  Uplink Inflight Pkts: 0                         Downlink Inflight Pkts: 0
  QFI Mismatch Uplink Pkts Drop: 0                QFI Mismatch Downlink Pkts Drop: 0

  Total QoS-Group(s) Active: 0
  DNS-to-EDNS Uplink Pkts: 0                     DNS-to-EDNS Uplink Bytes: 0
  EDNS Response Received: 0
```

```
  Flow information:
    Current Active Flows:
      TCP: 0
      UDP: 0
    Total Flows:
      TCP: 0
      UDP: 0
      FP:  0

Static & Predef Rule Match stats:
Rule Name            Pkts-Down  Bytes-Down Pkts-Up    Bytes-Up   Hits       Match-Bypassed
  FP-Down(Pkts/Bytes)  FP-Up(Pkts/Bytes)
-------------------- ---------- ---------- ---------- ---------- ---------- --------------
  -------------------  -----------------

Dynamic Rule Match stats:
PDR Id   Rule Name            Pkts-Down  Bytes-Down Pkts-Up    Bytes-Up   Hits
Match-Bypassed  FP-Down(Pkts/Bytes)  FP-Up(Pkts/Bytes)
-------- -------------------- ---------- ---------- ---------- ---------- ----------
--------------  -------------------  -----------------

Post-Processing Rule Match stats:
Rule Name            Pkts-Down  Bytes-Down Pkts-Up    Bytes-Up   Hits       Match-Bypassed
  FP-Down(Pkts/Bytes)  FP-Up(Pkts/Bytes)
-------------------- ---------- ---------- ---------- ---------- ---------- --------------
  -------------------  -----------------

Firewall-NAT Access Rule Match stats:
Firewall-Ruledef Name Pkts-Down Bytes-Down    Pkts-Up   Bytes-Up Dr-Pkts-Dn Dr-Byts-Dn
Dr-Pkts-Up Dr-Byts-Up       Hits
-------------------- ---------- ---------- ---------- ---------- ---------- ----------
---------- ---------- ----------

QoS-Group Statistics:
QGR Name             Pkts-Down  Bytes-Down Pkts-Up    Bytes-Up   Hits       Match-Bypassed
  FP-Down(Pkts/Bytes)  FP-Up(Pkts/Bytes)
-------------------- ---------- ---------- ---------- ---------- ---------- --------------
  -------------------  -----------------

  CIoT APN Rate Control:
    Allowed UL Limit:   0                        Allowed DL limit:    0
    Remaining UL Limit: 0                        Remaining DL limit:  0
    Allowed Time unit: unrestricted              Status Validity Time: N/A

  Local SEID      : [0x0004000000000002] 1125899906842626
  Remote SEID     : [0x0004000000000009] 1125899906842633
  State           : Connected
  Connect Time    : Thu Sep 15 20:51:03 2022
Idle time       : 00h00m31s
  Access Type: uplane-ipv4               Network Type: IP
  user-plane-service-name: UP21
  active-service-scheme-name:
  Callid: 00004e23
  Rulebase: cisco
  Interface Type: Sxb
  eMPS Session: No
  eMPS Session Priority: 0
  Card/Cpu: 1/0                          Sessmgr Instance: 1
  IP address: 11.11.12.1
  Next Hop Ip Address:
  Source context: EPC2-UP                Destination context: ISP1-UP
  PDN-Instance: starent.com
  User-plane-Sx-addr: 20.20.20.106
  Control-plane-Sx-addr: 20.20.20.101
```

```
 Number of associated PDRs : 0
 Number of associated ADC PDRs : 0
 Number of associated FARs : 0
 Number of associated QERs : 1
 Number of associated BARs : 0
 Number of associated URRs : 0
 Uplink APN AMBR (bps) : 200000       Downlink APN AMBR (bps) : 200000
 active input acl: IPV4ACL            active output acl: IPV4ACL
 active input ipv6 acl: n/a             active output ipv6 acl: n/a
 Bandwidth Policy: bw_policy1
 FW-and-NAT Policy: n/a
 FW-and-NAT Policy ID: n/a
 Firewall Policy IPv4: n/a
 Firewall Policy IPv6: n/a
 NAT Policy NAT44: n/a
 NAT Policy NAT64: n/a
 Converged Session: No                 Converged Peer Callid:      n/a
 Visited Call: No
 Subscriber Parameters:
 IMSI: 404005123456789
 IMEI: 112233445566778
 MSISDN: 9890098900
 Charging Characteristics: 0x3412
 Rat Type: 6
 MCC MNC: 123765
 SGSN Address: 20.20.20.61
 ULI: TAI=12546312ab00
 Congestion Level: 0
 GGSN Address: 20.20.20.101
 User-Name: 9890098900@starent.com
 Session-ID: 141414650000000A
 input pkts: 0                                 output pkts: 0
 input bytes: 0                                output bytes: 0
 input bytes dropped: 0                        output bytes dropped: 0
 input pkts dropped: 0                         output pkts dropped: 0
 CF Buffered Uplink Packets: 0                 CF Buffered Downlink Packets: 0

 CF Buffered Uplink Bytes: 0                   CF Buffered Downlink Bytes: 0
 Uplink Packets in Buffer: 0                   Downlink Packets in Buffer: 0
 Buff Over-limit Uplink Pkts: 0                Buff Over-limit Downlink Pkts: 0

 DDN buffered pkts : 0                         DDN buffered bytes : 0
 DDN buffer overflow drop pkts : 0             DDN buffer overflow drop bytes : 0

 pk rate from user(bps): 0                     pk rate to user(bps): 0
 ave rate from user(bps): 0                    ave rate to user(bps): 0
 sust rate from user(bps): 0                   sust rate to user(bps): 0
 pk rate from user(pps): 0                     pk rate to user(pps): 0
 ave rate from user(pps): 0                    ave rate to user(pps): 0
 sust rate from user(pps): 0                   sust rate to user(pps): 0
 ipv4 bad hdr: 0                               ipv4 ttl exceeded: 0
 ipv4 fragments sent: 0                        ipv4 could not fragment: 0
 ipv4 bad length trim: 0
 input pkts dropped (0 mbr): 0                 output pkts dropped (0 mbr): 0
 ipv4 input acl drop: 0                        ipv4 output acl drop: 0
 ipv6 input acl drop: 0                        ipv6 output acl drop: 0
 ip source violations: 0                       ipv4 output no-flow drop: 0
 ipv6 bad hdr: 0                               ipv6 bad length trim: 0
 ipv4 icmp packets dropped: 0
 APN AMBR Uplink Pkts Drop: 0                   APN AMBR Downlink Pkts Drop: 0
 APN AMBR Uplink Bytes Drop: 0                  APN AMBR Downlink Bytes Drop: 0

 APN AMBR Uplink Pkts IP pref lowered: 0          APN AMBR Downlink Pkts IP pref lowered:
0
```

```
  APN AMBR Uplink Bytes IP pref lowered: 0         APN AMBR Downlink Bytes IP pref lowered:
0
  ITC Uplink Pkts Drop: 0                           ITC Downlink Pkts Drop: 0
  ITC Uplink Bytes Drop: 0                          ITC Downlink Bytes Drop: 0
  ITC Uplink Pkts IP pref lowered: 0                ITC Downlink Pkts IP pref lowered: 0

  ITC Uplink Bytes IP pref lowered: 0               ITC Downlink Bytes IP pref lowered: 0

  ITC Terminated Flows:: 0                          ITC Redirected Flows: 0
  Flow Action Terminated Flows: 0                   Flow Action Redirected Flows: 0

  ToS marked Uplink Pkts: 0                         ToS marked Downlink Pkts: 0
  CC Dropped Uplink Pkts: 0                         CC Dropped Downlink Pkts: 0
  CC Dropped Uplink bytes: 0                        CC Dropped Downlink Bytes: 0
  Uplink Inflight Pkts: 0                           Downlink Inflight Pkts: 0
  QFI Mismatch Uplink Pkts Drop: 0                  QFI Mismatch Downlink Pkts Drop: 0

  Total QoS-Group(s) Active: 0
  DNS-to-EDNS Uplink Pkts: 0                      DNS-to-EDNS Uplink Bytes: 0
  EDNS Response Received: 0

  Flow information:
    Current Active Flows:
      TCP: 0
      UDP: 0
    Total Flows:
      TCP: 0
      UDP: 0
      FP:  0

Static & Predef Rule Match stats:
Rule Name            Pkts-Down  Bytes-Down Pkts-Up    Bytes-Up   Hits       Match-Bypassed
  FP-Down(Pkts/Bytes)  FP-Up(Pkts/Bytes)
-------------------- ---------- ---------- ---------- ---------- ---------- --------------
  -------------------  ------------------

Dynamic Rule Match stats:
PDR Id   Rule Name            Pkts-Down  Bytes-Down Pkts-Up    Bytes-Up   Hits
Match-Bypassed  FP-Down(Pkts/Bytes)  FP-Up(Pkts/Bytes)
-------- -------------------- ---------- ---------- ---------- ---------- ----------
--------------  -------------------  ------------------

Post-Processing Rule Match stats:
Rule Name            Pkts-Down  Bytes-Down Pkts-Up    Bytes-Up   Hits       Match-Bypassed
  FP-Down(Pkts/Bytes)  FP-Up(Pkts/Bytes)
-------------------- ---------- ---------- ---------- ---------- ---------- --------------
  -------------------  ------------------

Firewall-NAT Access Rule Match stats:
Firewall-Ruledef Name Pkts-Down Bytes-Down    Pkts-Up   Bytes-Up Dr-Pkts-Dn Dr-Byts-Dn
Dr-Pkts-Up Dr-Byts-Up       Hits
-------------------- ---------- ---------- ---------- ---------- ---------- ----------
---------- ---------- ----------

QoS-Group Statistics:
QGR Name             Pkts-Down  Bytes-Down Pkts-Up    Bytes-Up   Hits       Match-Bypassed
  FP-Down(Pkts/Bytes)  FP-Up(Pkts/Bytes)
-------------------- ---------- ---------- ---------- ---------- ---------- --------------
  -------------------  ------------------

  CIoT APN Rate Control:
    Allowed UL Limit:   0                         Allowed DL limit:    0
    Remaining UL Limit: 0                         Remaining DL limit:  0
    Allowed Time unit: unrestricted               Status Validity Time: N/A
```

```
Total subscribers matching specified criteria: 1
```

show subscribers user-plane-only callid *callid-value* rule-match-info { [full] all | rulename *rulename*}

Table 137: Feature History

Feature Name	Release Information	Description
Identify URR Associated With Matching Rule Using a Single Command	2025.01.0	This release introduces a consolidated command for UP to provide a comprehensive overview of per subscriber details, including: • Matched rules and their associated charging actions • The associated URR-ID This enhancement eliminates the need to use multiple commands across the UP or UP and CP to extract and correlate this information per subscriber. **Command Introduced**: **show subscribers user-plane-only callid** *callid* **rule-match-info { [full] all | rulename** *rulename* **}**. **Default Setting**: Enabled by default.

The **show subscribers user-plane-only callid** *callid* **rule-match-info { [full] all | rulename** *rulename* **}** command displays rules that are matched and their associated charging-actions and the associated URR-ID.

Note

- This command displays static, predefined, and dynamic rule matching statistics. It does not display postprocessing, firewall, or any other ruledefs.
- URR-ID corresponding information threshold, and quota needs to be obtained using the **show subscribers user-plane-only callid urr** *urr-ID* command.
- The URR-IDs that are displayed in the command may not be currently present in the output of the **show subscribers user-plane-only callid** *callid* **urr full all**. However, the URR-IDs are populated eventually.
- Executing this command with the **rulename** filter can be CPU intensive.
- Post recovery, when the packet count is lost, the rule-match stats are not recovered but the URR details are recovered.

The following example shows a sample output of the **show subscribers user-plane-only callid 00004e22 rule-match-info full all** command.

```
[local]host_name# show subscribers user-plane-only callid 00004e22 rule-match-info full all

STATIC & PREDEF RULE MATCH INFORMATION:

        RULEDEF NAME: catchall
        Charging-Action name: catchall

        Pkts-Down  Bytes-Down Pkts-Up    Bytes-Up   Hits       Match-Bypassed
FP-Down(Pkts/Bytes)  FP-Up(Pkts/Bytes)
        ---------- ---------- ---------- ---------- ---------- --------------
-------------------  ------------------
                 1         40          2         80          5              0          0/0
              0/0

        RULEDEF NAME: httpAnyMatch
        Charging-Action name: catchall

        Pkts-Down  Bytes-Down Pkts-Up    Bytes-Up   Hits       Match-Bypassed
FP-Down(Pkts/Bytes)  FP-Up(Pkts/Bytes)
        ---------- ---------- ---------- ---------- ---------- --------------
-------------------  ------------------
                 2        716          3        205          5              0          0/0
              0/0


        URR-id: 0x80000131

         URR-id: 0x80000137

DYNAMIC RULE MATCH INFORMATION:

        RULEDEF NAME: dynamic5

        PDR Id   Pkts-Down  Bytes-Down Pkts-Up    Bytes-Up   Hits       Match-Bypassed
FP-Down(Pkts/Bytes)  FP-Up(Pkts/Bytes)
        -------- --------   ---------- ---------- ---------- -------    ----------
--------------       -------------------
        0x0003            0          0          5        285          6              0
      0/0              0/0
        0x0004            3        756          0          0          4              0
```

```
      0/0              0/0

        Outer header removal:   GTP-U/UDP/IPv4
        Application ID:         N/A
        Associated FARID:       0x0003
         Destination Interface: Core
         Apply Action:          FORWARD
         Outer Header Creation:
         Remote TEID:           0x0
          Transport Level Marking:         N/A
          Transport Level Marking Options:
            Copy Inner:                    No
            Copy Outer:                    No
          Inner Packet Marking:            N/A

        Outer header removal:
        Application ID:         N/A
        Associated FARID:       0x0004
         Destination Interface: Access
         Apply Action:          FORWARD
         Outer Header Creation: GTP-U/UDP/IPv4
         Remote TEID:           0x60010f20
          Transport Level Marking:         N/A
          Transport Level Marking Options:
            Copy Inner:                    No
            Copy Outer:                    No
          Inner Packet Marking:            N/A

        URR-id: 0x00000007

Total subscribers matching specified criteria: 1
```

VPP Statistics

To determine if the flows are offloaded to VPP, check for Fastpath statistics in the output of the following CLI commands:

- **show user-plane-service statistics all**
- **show user-plane-service statistics analyzer name ip [verbose]**
- **show user-plane-service statistics analyzer name ipv6 [verbose]**
- **show user-plane-service statistics analyzer name tcp [verbose]**
- **show user-plane-service statistics analyzer name udp [verbose]**
- **show user-plane-service statistics analyzer name http [verbose]**
- **show user-plane-service statistics analyzer name rtp [verbose]**
- **show user-plane-service statistics analyzer name ftp [verbose]**

```
show user-plane-service statistics analyzer name ftp verbose

Wednesday July 02 17:47:13 IST 2025
   FTP Session Stats:
      Current Control Sessions:         0         Current Data Sessions:         0
      Total Control Sessions:           0         Total Data Sessions:           1
      Uplink Control Bytes:             0         Downlink Control Bytes:        0
      Uplink Control Pkts:              0         Downlink Control Pkts:         0
```

```
Uplink Data Bytes:              219         Downlink Data Bytes:         3249
Uplink Data Pkts:                 6         Downlink Data Pkts:             9
Uplink Error Bytes:               0         Downlink Error Bytes:           0
Uplink Error Pkts:                0         Downlink Error Pkts:            0
Request Succeed:                  0         Request Failed:                 0
Unknown Requests:                 0         Unknown Responses:              0
Uplink Bytes Retrans:             0         Downlink Bytes Retrans:         0
Uplink Pkts Retrans:              0         Downlink Pkts Retrans:          0
RETR commands:                    0         STOR commands:                  0
Unknown packets received:                                                   0
Data packet received without control connection:                           11
Invalid packets:                                                            0
Packets that could not be parsed:
```

- **show subscribers user-plane-only full callid** *call_id*
- **show subscribers user-plane-only callid** *callid_value* **drop-statistics**
- **show user-plane-service statistics drop-counter**
- **show subscribers user-plane-only callid** *callid_value* **flows full**
- **show subscribers user-plane-only callid** *callid_value* **flows flow-id** *flow_id*

show subscribers user-plane-only callid *callid_value* drop-statistics

The **show subscribers user-plane-only callid** *callid_value* **drop-statistics** command displays the packet drop statistics with the respective cause (drop reason) at instance-level and session-level.

The following is a sample output of this command:

```
[local]UPF1# show subscribers user-plane-only callid 00004e21 drop-statistics

         Callid: 00004e21
         Interface Type: <interface-type>

         Packet Drop Data Statistics:
         ----------------------------------------------------
         NAT packets processing failure
            NAT on demand handling:                  0
            ICMP Packet translation:                 0
         FIREWALL packets processing failure
            Policy not found:                        0
         No Matching GX rule found:                  0
         Flow apply action
            Discard:                                 0
            Readdress Failure:                       0
            Packet exceeds the MTU size:                 0
       Failure in processing FAR Buffer packets:     0
       FAR Apply Action Drop:                        0
       Traffic Steering Failure:                     0
       QER Gate Status Closed:                       0
       Content-filtering Discard Action:             0
       IP Header Validation Failed:                  0
       ADF level failure
           DL TFT mismatch:                          0
       URL Blacklisting Discard Action:              0
       QGR Flow Action SGQ Discard:                  0
       QGR Policer Drops:                            0

Total subscribers matching specified criteria: 1
[local]UPF1#
```

show user-plane-service statistics drop-counter

The **show user-plane-service statistics drop-counter** command displays the packet drop statistics with the respective cause (drop reason) at instance-level and session-level.

The following is a sample output of this command:

```
[local]UPF1# show user-plane-service statistics drop-counter

        Packet Drop Data Statistics:
        ----------------------------------------------------
        NAT packets processing failure
           NAT on demand handling:                      0
           IP allocation is in progress:                0
           ICMP Packet translation:                     0
        FIREWALL packets processing failure
           Policy not found:                            0
        No Matching GX rule found:                      0
        Flow apply action
           Discard:                                     0
           Readdress Failure:                           0
           Redirect-URL:                                0
        Packet exceeds the MTU size:                    0
        Failure in processing FAR Buffer packets:       0
        FAR Apply Action Drop:                          0
        Traffic Steering Failure:                       0
        QER Gate Status Closed:                         0
        Content-filtering Discard Action:               0
        IP Header Validation Failed:                    0
        ADF level failure
          UL TEID/QFI key mismatch:                     0
          DL TFT mismatch:                              0
          DL QFI mismatch:                              0
        URL Blacklisting Discard Action:                0
        DDN buffer overflow drop packets:               0
        APN AMBR Packets Drop:                          0
        ITC Packets Drop:                               0
        ACL Drop:                                       0
        CC Dropped Packets:                             0
        FastPath Misc Drops
            Overload Protection:                        0
            Invalid Client:                             0
            Stream ID 0:                                0
            Invalid Stream ID:                          0
```

NOTES:

- For the **show user-plane-service statistics drop-counter** CLI command, the counters for the number of packets dropped due to various reasons are not recovered post sessmgr recovery or ICSR (Inter-Chassis Session Recovery).
- For the **show subscribers user-plane-only callid** *callid_value* **drop-statistics** CLI command, the counters for the number of packets dropped due to various reasons are recovered post sessmgr recovery or ICSR.
- The commands do not duplicate the counters that are captured at the session-level and instance-level.
- The packets dropped due to "Quota Exhaust FAR Apply Action Drop" are accumulated under "CC Dropped Packets" and not under the "FAR Apply Action Drop" counter.
- The packets dropped under "Failure in processing FAR buffered packets" can be due to "no rule match", "teid not found", or other reasons. The respective counter are incremented for such packets.

- The packets dropped due to the "redirect-url" flow action are accumulated under "Redirect-URL" in the **show user-plane-service statistics drop-counter** CLI command.

show subscribers user-plane-only callid *callid_value* flows full

The **show subscribers user-plane-only callid** *callid_value* **flows full** command displays the detailed information at the flow level which are required for debugging purposes.

The following is a sample output of this command:

```
[local]GVK-HUPF# show subscribers user-plane-only callid 017dc662 flows full
Tuesday November 22 18:35:35 UTC 2022

Callid: 017dc662
Interface Type: N4
IP address: n/a

Flow ID: 6:2
Uplink pkts: 0 Downlink pkts: 1
Uplink bytes: 0 Downlink bytes: 1040
Fast Path Info:
Uplink pkts: 0 Downlink pkts: 0
Uplink bytes: 0 Downlink bytes: 0
Total pkts: 0
Uplink Dropped pkts: 0 Downlink Dropped pkts: 4
Uplink Stream ID: 0x0 Downlink Stream ID: 0x1000007
Uplink Stream State: n/a Downlink Stream State: Passive
Client-ID : C6
UE IP address: 12.1.0.3 UE Port: 1001
Server IP address: 100.100.100.10 Server Port: 1001
Protocol: TCP
Service Chain Name: NA
Uplink Sfp Id: NA
Downlink Sfp Id: NA

Flow ID: 6:3
Uplink pkts: 5 Downlink pkts: 2
Uplink bytes: 200 Downlink bytes: 2080
Fast Path Info:
Uplink pkts: 0 Downlink pkts: 3
Uplink bytes: 0 Downlink bytes: 3120
Total pkts: 3
Uplink Dropped pkts: 0 Downlink Dropped pkts: 0
Uplink Stream ID: 0x0 Downlink Stream ID: 0x2000008
Uplink Stream State: n/a Downlink Stream State: Active
Client-ID : C6
UE IP address: 12.1.0.3 UE Port: 1000
Server IP address: 100.100.100.10 Server Port: 1000
Protocol: TCP
Service Chain Name: NA
Uplink Sfp Id: NA
Downlink Sfp Id: NA

Total Number of Active flows : 2
```

show subscribers user-plane-only callid *callid_value* flows flow-id *flow_id*

The **show subscribers user-plane-only callid** *callid_value* **flows flow-id** *flow_id* command displays the detailed information at the flow level which are required for debugging purposes.

The following is a sample output of this command:

```
# show sub user-plane-only callid 00004e22 flows flow-id 1:4

  Callid: 00004e22
  Interface Type: Sxb
  IP address: n/a

  Flow ID: 1:4
  Uplink pkts: 17                                 Downlink pkts: 0
  Uplink bytes: 1026                              Downlink bytes: 0
  Fast Path Info:
   Uplink pkts: 0                                 Downlink pkts: 0
   Uplink bytes: 0                                Downlink bytes: 0
   Total pkts: 0
   Uplink Dropped pkts: 0                         Downlink Dropped pkts: 0
   Uplink Stream ID: 0x0000008                    Downlink Stream ID: 0x0
   Uplink Stream State: Passive                   Downlink Stream State: n/a
   Client-ID : C2
  UE IP address: 12.2.0.13                        UE Port: 1007
  Server IP address: 100.100.100.10               Server Port: 1005
  Protocol: TCP
  Service Chain Name: NA
  Uplink Sfp Id: NA
  Downlink Sfp Id: NA

Total Number of Active flows : 1
```

SNMP Support

The system uses the Simple Network Management Protocol (SNMP) to send traps or events to the EMS server or an alarm server on the network. You must configure SNMP settings to communicate with those devices.

The *SNMP MIB Reference* describes the MIBs and SNMP traps that are supported by UPF and StarOS.

The following SNMP traps are available in support of their respective feature or functionality:

N4 Session/Node Level Reporting Procedure

The following traps are available to track status and conditions GTP-U path failure:

- EGTPUPathFailure: This trap is generated when no response is received for GTP-U ECHO requests and data path failure is detected toward peer EPC Node.
- EGTPUPathFailureClear: This trap is generated when the data path toward the peer node is available.

UP Session Recovery

The following traps are available after session recovery in the User Plane node:

- ManagerFailure: This trap is generated when there is failure in the Software manager.
- TaskFailed: This trap is generated when a noncritical task has failed and the appropriate recovery steps begin.
- TaskRestart: This trap is generated when a noncritical task has restarted after an earlier failure.
- SessMgrRecoveryComplete: This trap is generated when Session Manager recovery completes. This is typically caused by the failure of Session Manager task and successful completion of recovery.
- ManagerRestart: This trap is generated when the identified manager task has been restarted.

Sx Association

The following traps are available to track the status of an Sx Association:

- SxPeerAssociated: This trap is triggered when an Sx association is detected.
- SxPeerAssociationRelease: This trap is triggered when an Sx association release is detected.

URL Blockedlisting

The following SNMP trap are available in support of URL Blockedlisting feature:

- BLDBError: Specifies the blockedlisting OPTBLDB file error that is displayed with an error code.
- BLDBErrorClear: Specifies the blockedlisting OPTBLDB file error removed.
- BLDBUpgradeError: Specifies the blockedlisting OPTBLDB file error displayed with an error code.
- BLDBUpgradeErrorClear: Specifies the blockedlisting OPTBLDB file error removed.

Enabling SNMP Traps

Use the following configuration to enable an SNMP trap.

```
configure
  snmp trap enable trap_name
  end
```

For supplemental information about SNMP Support, see *Management Settings* chapter in the *ASR 5500 System Administration Guide*.

Troubleshooting UPF Features

N4 or Datapath

The following CLI commands are available for troubleshooting N4 or datapath related issues:

- **show gtpu statistics**
- **show user-plane-service { all | bandwidth-policy | charging-action | edr-format | group-of-ruledefs | gtpp-group | name | pdn-instance | rulebase | ruledef | statistics | xheader-format }**

  NOTES:

  - **all**: Displays all User Plane services.
  - **bandwidth-policy**: Displays information for bandwidth-policy in User Plane service.
  - **charging-action**: Displays information for Charging actions in User Plane service.
  - **edr-format**: Displays information for EDR format in user Plane service.
  - **group-of-ruledefs**: Displays information on Group of Ruledefs configured in User Plane service.
  - **gtpp-group**: Displays information for bandwidth policy in User Plane service.
  - **name**: Displays information for specific User Plane service name.

- **pdn-instance**: Displays information for PDN instance.
- **rulebase**: Displays information for rulebase in User Plane service.
- **ruledef**: Displays information for ruledef in User Plane service.
- **statistics**: Displays node-level statistics for User Plane.

  Additionally, you can also use: **show user-plane-service statistics { all | analyzer | charging-action | fapi | rulebase | tethering-detection }**
- **xheader-format**: Displays information for X-Header format in User Plane service.

- **show user-plane-service content-filtering category policy-id (all | id** *id_value* **}**
  - **content-filtering**: Displays content filtering information.
  - **category**: Displays content filtering category information.
  - **policy-id**: Displays content filtering category Policy-ID and its definition.
  - **all**: Displays definitions of all content filtering category policies.
  - **id** *id_value*: Displays content filtering category definition of a particular Policy-ID. *id_value* is an integer ranging from 1 through 4,294,967,295.

- **show sx-service { all | name | statistics }**

  NOTES:
  - **all**: Displays all Sx Services.
  - **name**: Displays information for specific Sx Service name.
  - **statistics**: Displays the total of collected information for specific protocol since last restart or clear command.

Content Filtering

Use the following CLI command for troubleshooting CF related issues:

In releases prior to 2022.01.0:

```
show user-plane-service inline-services { content-filtering | info |
url-blacklisting }
```

From 2022.01.0 and later releases:

```
show user-plane-service inline-services { content-filtering | info |
url-blockedlisting }
```

NOTES:

- **content-filtering**: Displays content filtering information.
- **info**: Displays information of inline services.
- **url-blockedlisting**: Displays URL Blockedlisting parameters in User Plane service.

URL Blacklisting

Use the following CLI command for troubleshooting URL Blacklisting related issues: **show user-plane-service url-blacklisting database { all | debug-only | facility | url }**

NOTES:

- **all**: Displays all URL Blacklisting database configurations.
- **debug-only**: Displays the URL Blacklisting static database debug information.
- **facility**: Displays URL Blacklisting database configuration per facility.
- **url**: Displays particular database information for URL Blacklisting.

Configure Stats Profile for QCI Drop Statistics

Feature name	Release	Feature description
Per-QCI Packet Drop Counters	2026.03.0	Added support for per-APN, per-QCI packet drop statistics in UPF. The feature introduces configurable stats profiles, enhanced show commands, and reporting of packet drops across 19 consolidated drop categories, enabling improved visibility and troubleshooting of traffic-specific packet loss, including VoLTE and other QoS-sensitive services.

Enable granular tracking and reporting of packet drop counters segmented by QCI value and specific drop reasons.

Note One statistics profile can be associated across muliple APNs.

Follow these steps to configure per-QCI drop statistics:

Procedure

Step 1 Create a statistics profile:

```
Copy Code
configure
  stats-profile <profile_name>
    qci [ all | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | non-std { gbr | non-gbr } ] arp <value>
  end
```

Step 2 Ensure the procedure reflects that you are associating the profile with the APN.

```
bash
Copy Code
configure
  context <ctx_name>
    apn <apn_name>
      stats-profile <profile_name>
    end
```

NSO configuration considerations

When configuring Per-APN Per-QCI Drop Statistics through NSO using the UPF NED, the `arp` parameter in the `stats-profile` configuration is not supported on the UPF. As a result, the UPF returns the following warning:

```
Warning: ARP filtering is not supported on UPF.
ARP values will be ignored and the entry will be tracked per-QCI only.
```

By default, the NSO UPF NED treats this warning as a configuration failure and aborts the commit operation. To commit the configuration successfully, configure the warning as an accepted configuration warning in the StarOS NED. Use the following command:

```
devices global-settings ned-settings cisco-staros behaviour \
config-warning-ignore \
"Warning: ARP filtering is not supported on UPF. ARP values will be ignored and the entry
will be tracked per-QCI only"
```

After configuring the warning as an accepted configuration warning, the Per-APN Per-QCI Drop Statistics configuration can be committed successfully through NSO.

Note The `arp` parameter is ignored on the UPF, and packet drop statistics are maintained on a per-QCI basis only.

Per-QCI Drop statistics commands

Use the following command to display the granular distribution of packet drop counters per QCI.

```
show user-plane-service stats-profile name <name>
```

Example output

The following example displays per-QCI packet drop statistics for the statistics profile:

```
show user-plane-service stats-profile name qci_selfprot_5g_prof
```

```
 ------------------------------------------------------------------------------
 QCI 5   (Non-GBR / IMS Signaling)      [Profile-enabled]
   Packet Drop Reason                                Uplink      Downlink
   ---------------------------------------- ------------- -------------
   Self Protection                                      1             1
   Others                                               1             1
   ---------------------------------------- ------------- -------------
   Total                                                2             2
 ------------------------------------------------------------------------------
```

Display Configured Statistics Profiles

Use the following command to display all configured statistics profiles.

```
show stats-profile all
```

Example output

```
show stats-profile all
Statistics Profiles
--------------------------------
Stats Profile Name : sp1
qci all arp all

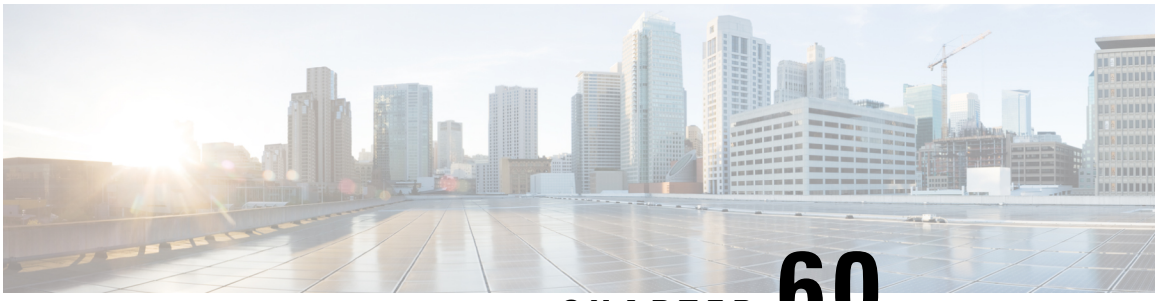
Stats Profile Name : sp2
qci 1 arp 1
```

Note The packet drop statistics displayed by the `show user-plane-service stats-profile` command include only drops accounted for by the user-plane service. Packets dropped at the VPP level are not included in these statistics.

Table 138: Drop reason description

Drop Reason	Description
Flow Action Discard	Packets dropped due to flow action configured as discard.
DDN Buffer Overflow	Packets dropped because the DDN buffer exceeded its capacity.
FAR Apply Action Drop	Packets dropped due to FAR action set to drop.
ITC Packets Drop	Packets dropped by ITC processing.
APN AMBR Packets Drop	Packets dropped due to APN AMBR rate enforcement.
CC Dropped Packets	Packets dropped because of charging or credit control restrictions.
QER Gate Status Closed	Packets dropped when the associated QER gate is closed.
NAT Packets Processing Failure	Packets dropped due to NAT processing failures.
FAR Buffer Pkt Processing Failure	Packets dropped because of FAR buffer processing failures.
Packet Exceeds MTU Size	Packets dropped because packet size exceeds the configured MTU.
IP Header Validation Failed	Packets dropped due to invalid or malformed IP headers.
ACL Drop	Packets dropped by Access Control List (ACL) rules.
SGACL Packet Drops	Packets dropped by Security Group ACL (SGACL) policies.
FastPath Misc Drops	Packets dropped for miscellaneous FastPath reasons.
ADF Level Failure	Packets dropped due to ADF processing failures.
QGR Policer Drop	Packets dropped by QGR policing mechanisms.
Overload Protection	Packets dropped to protect the system during overload conditions.
Self Protection	Packets dropped by self-protection mechanisms.

Drop Reason	Description
Others	Packets dropped for reasons not covered by the above categories.



CHAPTER 60

Sample UPF Configuration

- Sample Configuration, on page 647

Sample Configuration

The following is only a sample configuration file provided solely for your reference. You must create and modify your own configuration file according to the specific needs of your deployment.

```
----- snip -----
  active-charging service acs
    bandwidth-policy BWP
      flow limit-for-bandwidth id 1 group-id 2
      flow limit-for-bandwidth id 2 group-id 3
      flow limit-for-bandwidth id 100 group-id 100
     group-id 2 direction uplink peak-data-rate 256000 peak-burst-size 1000 violate-action
 discard
     group-id 2 direction downlink peak-data-rate 200000 peak-burst-size 1000 violate-action
 discard
     group-id 3 direction downlink peak-data-rate 256000 peak-burst-size 1000 violate-action
 discard
     group-id 4 direction uplink peak-data-rate 300000 peak-burst-size 1200 violate-action
 lower-ip-precedence
      group-id 100 direction downlink peak-data-rate-kbps 4294967295 peak-burst-size
4294967295 violate-action discard

    ruledef L3_SERVER
      ip server-ip-address = 209.165.202.150/27
      tcp either-port = 80
    #exit
    ruledef L4_PORT
      tcp either-port = 80
      udp either-port = 80
      multi-line-or all-lines
    #exit
    ruledef L7_HTTP
      http host contains 209.165.202.150
      multi-line-or all-lines
    #exit
    ruledef http-pkts
      http any-match = TRUE
    #exit
    ruledef http-port
      tcp either-port = 80
      rule-application routing
    #exit
```

```
  ruledef ip-any-rule
    ip any-match = TRUE
  #exit
  urr-list urrs
    rating-group 10 urr-id 5
  #exit
  charging-action starent
    content-id 10
    billing-action egcdr
  #exit
  rulebase default
  #exit
  credit-control group default
    pending-traffic-treatment noquota buffer
    pending-traffic-treatment quota-exhausted pass
    usage-reporting quotas-to-report based-on-grant

  rulebase starent
    billing-records egcdr
    dynamic-rule order first-if-tied
    action priority 5 ruledef http-pkts charging-action standard
    action priority 10 ruledef L7_HTTP charging-action starent
    action priority 20 ruledef L4_PORT charging-action starent
    action priority 100 ruledef L3_SERVER charging-action starent
    action priority 10000 ruledef ip-any-rule charging-action starent
    route priority 1 ruledef http-port analyzer http
    egcdr threshold interval 1000
    bandwidth default-policy BWP
  #exit
  traffic-optimization-policy default
  #exit
#exit
context ingress
  interface N3_interface
    ip address 209.165.201.4 209.165.201.5
    ipv6 address abc0:0:0:cb::1/64 secondary
  #exit
  interface N3_interface_LOGICAL loopback
    ip address 209.165.201.4 209.165.201.5
  #exit
  interface N3_interface_LOGICAL2 loopback
    ip address 209.165.201.4 209.165.201.5
  #exit
  interface N4U_interface
    ip address 209.165.201.4 209.165.201.5
    ipv6 address abc0:0:0:cd::1/64 secondary
    ipv6 address abc0:0:0:ca::1/64 secondary
  #exit
  interface N4U_interface_LOGICAL loopback
    ip address 209.165.201.4 209.165.201.5
  #exit
 interface N4_interface
    ip address 209.165.201.4 209.165.201.5
    ipv6 address abc0:0:0:cc::1/64 secondary
  #exit
  interface N4_interface_LOGICAL loopback
    ip address 209.165.201.4 209.165.201.5
  #exit
  subscriber default
  exit
  aaa group default
  #exit
  gtpp group default
  #exit
```

```
  gtpu-service N3-GNB1
    bind ipv4-address 209.165.200.225
  exit
  gtpu-service N3-GNB2
    bind ipv4-address 209.165.201.4
  exit
  gtpu-service control_gtpu
    bind ipv4-address 209.165.201.4
  exit
  sx-service N4
    instance-type userplane
    bind ipv4-address 209.165.201.4
    sx-protocol heartbeat interval 3600
    sx-protocol heartbeat max-retransmissions 1
    sx-protocol association reattempt-timeout 30
  exit
  user-plane-service user-plane-service
    associate gtpu-service N3-GNB1 upf-ingress
    associate gtpu-service control_gtpu cp-tunnel
    associate sx-service N4
    associate fast-path service
    associate control-plane-group default
    load-control capacity 900
  exit
  user-plane-service user-plane-service1
  exit
  ip route 209.165.201.4 209.165.201.5 209.165.201.6 N4_interface
  ip route 209.165.201.4 209.165.201.5 209.165.201.6 N4_interface
  ip route 209.165.201.4 209.165.201.5 209.165.201.6 N3_interface
  ip route 209.165.201.4 209.165.201.5 209.165.201.6 N3_interface
  ip route 209.165.201.4 209.165.201.5 209.165.201.6 N4_interface
  ip route 209.165.201.4 209.165.201.5 209.165.201.6 N3_interface
  ip route 209.165.201.4 209.165.201.5 209.165.201.6 N4_interface
  ip route 209.165.201.4 209.165.201.5 209.165.201.6 N3_interface
#exit
context egress
  interface N6_interface
    ip address 209.165.201.4 209.165.201.5
    ipv6 address abc0:0:0:cf::1/64 secondary
  #exit
  subscriber default
  exit
  apn starent.com
    pdp-type ipv4 ipv6
    selection-mode subscribed sent-by-ms chosen-by-sgsn
    gtpp group default accounting-context egress
    ip context-name egress
    active-charging rulebase starent
  exit
  aaa group default
  #exit
  gtpp group default
    gtpp attribute local-record-sequence-number
    gtpp dictionary custom24
    gtpp egcdr service-data-flow threshold interval 60
    gtpp egcdr service-data-flow threshold volume downlink 13000
    gtpp egcdr service-data-flow threshold volume uplink 17000
    gtpp egcdr service-data-flow threshold volume total 22222
  #exit
  ipv6 route 2:2:2:2::/64 next-hop abc0::ab:1c:2ff:def9:1ad interface N6_interface
  ip route 209.165.201.4 209.165.201.5 209.165.201.6 N6_interface
  ip route 209.165.201.4 209.165.201.5 209.165.201.6 N6_interface
  ipv6 route 2:2:2:2::/64 next-hop abc0::ab:1c:2ff:def9:1ab interface N6_interface
  ip route 209.165.201.4 209.165.201.5 209.165.201.6 N6_interface
```

```
  #exit
  control-plane-group default
    sx-association initiated-by-cp
    peer-node-id ipv4-address 209.165.200.225 interface n4
  #exit
  user-plane-group default
  #exit
  port ethernet 1/11
    no shutdown
    vlan 203
      no shutdown
      bind interface N4U_interface ingress
    #exit
    vlan 204
      no shutdown
      bind interface N4_interface ingress
    #exit
    vlan 205
      no shutdown
      bind interface N3_interface ingress
    #exit
    vlan 206
      no shutdown
    #exit
    vlan 207
      no shutdown
      bind interface N6_interface egress
    #exit
  #exit
end
```