

BGP Command Reference for Cisco NCS 5500 Series, Cisco NCS 540 Series, and Cisco NCS 560 Series Routers

First Published: 2015-12-23

Last Modified: 2026-09-11

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)

© 2015–2026 Cisco Systems, Inc. All rights reserved.



CONTENTS

PREFACE

Preface xi

Changes to This Document xi

Communications, Services, and Additional Information xii

CHAPTER 1

BGP Commands 1

accept-own 8

additional-paths advertise-limit 9

additional-paths install backup 11

additional-paths receive 12

additional-paths selection 14

additional-paths send 16

address-family (BGP) 18

advertise best-external 22

advertise permanent-network 23

advertisement-interval 24

advertise vpn-imported label-mode per-vrf 25

af-group 26

aggregate-address 28

aigp 30

aigp send-cost-community 31

allocate-label 33

allowas-in 35

allowconfedas-in 36

as-format 37

as-override 38

as-path-loopcheck out disable 39

attribute large-community 40
attribute-filter group 41
bfd (BGP) 42
bestpath as-path multipath-relax (BGP IPv4 RT-filter address family) 47
bgp as-path-loopcheck 48
bgp attribute-download 49
bgp auto-policy-soft-reset disable 51
bgp bestpath origin-as allow invalid 52
bgp bestpath origin-as use validity 53
bgp bestpath as-path ignore 54
bgp bestpath compare-routerid 55
bgp bestpath cost-community ignore 56
bgp bestpath med always 57
bgp bestpath med confed 59
bgp bestpath med missing-as-worst 61
bgp client-to-client reflection disable 62
bgp cluster-id 64
bgp confederation identifier 66
bgp confederation peers 68
bgp dampening 70
bgp default local-preference 72
bgp enforce-first-as disable 73
bgp fast-external-fallover disable 74
bgp graceful-restart 75
bgp graceful-restart graceful-reset 77
bgp graceful-restart purge-time 78
bgp graceful-restart restart-time 79
bgp graceful-restart stalepath-time 80
bgp import-delay 81
bgp label-delay 82
bgp lpts-secure-binding 83
bgp log neighbor changes disable 84
bgp origin-as validation enable 85
bgp origin-as validation signal ibgp 86

bgp origin-as validation time 87
bgp maximum neighbor 89
bgp multipath as-path 90
bgp policy accounting 91
bgp policy propagation input flow-tag 93
bgp redistribute-internal 94
bgp router-id 95
bgp scan-time 96
bgp unsafe-ebgp-policy 98
bgp update-delay 99
bgp write-limit 101
bmp-activate 103
bmp server 104
bind-source (rpki-server) 107
capability additional-paths receive 108
capability additional-paths send 109
capability orf prefix 110
capability suppress 4-byte-as 112
cef load-balancing recursive oor mode dampening-and-dlb 115
clear bgp 116
clear bgp dampening 118
clear bgp external 120
clear bgp flap-statistics 122
clear bgp nexthop performance-statistics 124
clear bgp nexthop registration 126
clear bgp peer-drops 128
clear bgp performance-statistics 129
clear bgp self-originated 130
clear bgp shutdown 132
clear bgp soft 134
clear bgp vrf rpki validation 137
cluster-id allow-equal 138
default-information originate (BGP) 139
default-martian-check disable 140

default-metric (BGP) 141
default-originate 142
default-prefix suppress (BGP) 144
description (BGP) 145
distance bgp 146
distribute bgp-ls (ISIS) 148
distribute bgp-ls (OSPF) 149
domain-distinguisher 150
dmz-link-bandwidth 151
dscp (BGP) 152
dynamic-med interval 154
ebgp-multihop 155
export route-policy 157
export route-target 158
hw-module fib bgp-mp-pic-auto-protect-enable 160
hw-module fib dlb level-1 enable 161
hw-module profile stats j2-dynamic-stats 163
hw-module fib bgppa stats-mode 164
hw-module fib urpf disable 166
ibgp policy out enforce-modifications 167
import route-policy 168
import route-target 169
ignore-connected-check 171
graceful-maintenance 172
hw-module profile flowspec 174
hw-module fib bgppa stats-mode 175
keychain (BGP) 177
keychain-disable 179
keychain inheritance-disable 181
label-allocation-mode 182
label mode 184
local-as 187
long-lived-graceful-restart 189
match flow-tag 191

maximum-paths (BGP) 193
maximum-paths external-rtfilter (BGP IPv4 RT-filter address family) 195
maximum-prefix (BGP) 196
maximum-prefix-restart-time (BGP) 199
mpls activate (BGP) 200
mvpn 203
neighbor (BGP) 204
neighbor-group 206
neighbor internal-vpn-client 208
network (BGP) 209
network backdoor 211
next-hop-self 213
next-hop-unchanged 215
nexthop resolution prefix-length minimum 217
nexthop route-policy 218
nexthop trigger-delay 220
orf 222
password (BGP) 223
password (rpki-server) 225
password-disable 226
permanent-network 227
precedence 228
preference (rpki-server) 230
prefix-ecmp-delay 231
purge-time (rpki-server) 232
rd 233
receive-buffer-size 235
redistribute (BGP) 237
refresh-time (rpki-server) 241
remote-as (BGP) 242
remote-bgp-label-mode-per-prefix 244
remove-private-as 246
replace as-path all 248
response-time (rpki-server) 250

retain local-label 251
retain route-target 252
route-policy (BGP) 254
router bgp 256
rpki server 257
selective-vrf-download disable 258
send-buffer-size 259
send-community-ebgp 261
send-extended-community-ebgp 263
session-group 265
session-open-mode 267
set flow-tag 269
set path-color external-reach 270
show bgp 271
show bgp bmp 282
show bgp update out 284
show bgp update in error process 288
show bgp update out filter-group 289
show bgp update out process 290
show bgp update out sub-group 292
show bgp update out update-group 293
show bgp vrf update in error 295
show bgp advertised 296
show bgp af-group 303
show bgp attribute-key 306
show bgp cidr-only 310
show bgp community 313
show bgp convergence 318
show bgp dampened-paths 321
show bgp flap-statistics 324
show bgp inconsistent-as 329
show bgp labels 333
show bgp large-community 336
show bgp l2vpn 339

show bgp l2vpn vpls 343
show bgp neighbor-group 349
show bgp neighbors 352
show bgp nexthops 379
show bgp paths 388
show bgp policy 390
show bgp process 397
show bgp regexp 417
show bgp route-policy 421
show bgp ipv4 rt-filter neighbor 425
show bgp session-group 426
show bgp sessions 429
show bgp summary 433
show bgp truncated-communities 437
show bgp update-group 441
show bgp vrf 447
show protocols (BGP) 450
show svd role 452
show svd state 453
shutdown (BGP) 454
shutdown (rpki-server) 456
signalling disable 457
site-of-origin (BGP) 458
slow peer (BGP router global configuration) 460
slow peer (BGP neighbor address-family configuration) 462
socket receive-buffer-size 465
socket send-buffer-size 467
soft-reconfiguration inbound 469
speaker-id 472
stats resource reassign 473
svd platform enable 478
table-policy 480
timers (BGP) 482
timers bgp 484

transport (rpki-server) 486
update limit 487
update limit address-family 488
update limit sub-group 489
update in error-handling basic disable 490
update in error-handling extended 491
update in error-handling treat-as-withdraw (BGP) 492
update in labeled-unicast equivalent 493
update out logging 494
update-source 495
update wait-install 497
update wait-install delay startup 498
use 499
username (rpki-server) 503
vrf (BGP) 504
weight 505



Preface

The *BGP Command Reference for Cisco NCS 5500 Series Routers* preface contains these sections:

- Changes to This Document, on page xi
- Communications, Services, and Additional Information, on page xii

Changes to This Document

This table lists the technical changes made to this document since it was first printed.

Date	Change Summary
December 2015	Initial release of this document.
November 2016	Republished with documentation updates for Release 6.1.2 features.
July 2017	Republished with documentation updates for Release 6.2.2 features.
September 2017	Republished with documentation updates for Release 6.3.1 features.
July 2018	Republished with documentation updates for Release 6.5.1 features.
June 2024	Republished with documentation updates for Release 24.2.1 features.
September 2025	Republished with documentation updates for Release 25.3.1 features.
December 2025	Republished with documentation updates for Release 25.4.1 features.
February 2026	Republished with documentation updates for Release 26.1.1 features.
June 2026	Republished with documentation updates for Release 26.2.1 features.

Communications, Services, and Additional Information

- To receive timely, relevant information from Cisco, sign up at Cisco Profile Manager.
- To get the business results you're looking for with the technologies that matter, visit Cisco Services.
- To submit a service request, visit Cisco Support.
- To discover and browse secure, validated enterprise-class apps, products, solutions and services, visit Cisco Marketplace.
- To obtain general networking, training, and certification titles, visit Cisco Press.
- To find warranty information for a specific product or product family, access Cisco Warranty Finder.

Cisco Bug Search Tool

Cisco Bug Search Tool (BST) is a web-based tool that acts as a gateway to the Cisco bug tracking system that maintains a comprehensive list of defects and vulnerabilities in Cisco products and software. BST provides you with detailed defect information about your products and software.



BGP Commands

Note All commands applicable for the Cisco NCS 5500 Series Router are also supported on the Cisco NCS 540 Series Router that is introduced from Cisco IOS XR Release 6.3.2. References to earlier releases in Command History tables apply to only the Cisco NCS 5500 Series Router.

Note

- Starting with Cisco IOS XR Release 6.6.25, all commands applicable for the Cisco NCS 5500 Series Router are also supported on the Cisco NCS 560 Series Routers.
- Starting with Cisco IOS XR Release 6.3.2, all commands applicable for the Cisco NCS 5500 Series Router are also supported on the Cisco NCS 540 Series Router.
- References to releases before Cisco IOS XR Release 6.3.2 apply to only the Cisco NCS 5500 Series Router.
- Cisco IOS XR Software Release 7.0.1 specific updates are not applicable for the following variants of Cisco NCS 540 Series Routers:
  - N540-28Z4C-SYS-A
  - N540-28Z4C-SYS-D
  - N540X-16Z4G8Q2C-A
  - N540X-16Z4G8Q2C-D
  - N540X-16Z8Q2C-D
  - N540-12Z20G-SYS-A
  - N540-12Z20G-SYS-D
  - N540X-12Z16G-SYS-A
  - N540X-12Z16G-SYS-D

This chapter describes the commands used to configure and monitor Border Gateway Protocol (BGP) on Cisco NCS 5500 Series Routers using Cisco IOS XR software. The commands in this module configure IP Version 4 (IPv4), IP Version 6 (IPv6), Virtual Private Network Version 4 (VPNv4) routing sessions.

For detailed information about BGP concepts, configuration tasks, and examples, see the Implementing BGP chapter in the *BGP Configuration Guide for Cisco NCS 5500 Series Routers*.

- accept-own, on page 8
- additional-paths advertise-limit, on page 9
- additional-paths install backup, on page 11
- additional-paths receive, on page 12
- additional-paths selection, on page 14
- additional-paths send, on page 16
- address-family (BGP), on page 18
- advertise best-external, on page 22
- advertise permanent-network, on page 23
- advertisement-interval, on page 24
- advertise vpn-imported label-mode per-vrf, on page 25
- af-group, on page 26
- aggregate-address, on page 28
- aigp, on page 30
- aigp send-cost-community, on page 31
- allocate-label, on page 33
- allowas-in, on page 35
- allowconfedas-in, on page 36
- as-format, on page 37
- as-override, on page 38
- as-path-loopcheck out disable, on page 39
- attribute large-community, on page 40
- attribute-filter group, on page 41
- bfd (BGP), on page 42
- bestpath as-path multipath-relax (BGP IPv4 RT-filter address family), on page 47
- bgp as-path-loopcheck, on page 48
- bgp attribute-download, on page 49
- bgp auto-policy-soft-reset disable, on page 51
- bgp bestpath origin-as allow invalid, on page 52
- bgp bestpath origin-as use validity, on page 53
- bgp bestpath as-path ignore, on page 54
- bgp bestpath compare-routerid, on page 55
- bgp bestpath cost-community ignore, on page 56
- bgp bestpath med always, on page 57
- bgp bestpath med confed, on page 59
- bgp bestpath med missing-as-worst, on page 61
- bgp client-to-client reflection disable, on page 62
- bgp cluster-id, on page 64
- bgp confederation identifier, on page 66
- bgp confederation peers, on page 68

- bgp dampening, on page 70
- bgp default local-preference, on page 72
- bgp enforce-first-as disable, on page 73
- bgp fast-external-fallover disable, on page 74
- bgp graceful-restart, on page 75
- bgp graceful-restart graceful-reset, on page 77
- bgp graceful-restart purge-time, on page 78
- bgp graceful-restart restart-time, on page 79
- bgp graceful-restart stalepath-time, on page 80
- bgp import-delay, on page 81
- bgp label-delay, on page 82
- bgp lpts-secure-binding, on page 83
- bgp log neighbor changes disable, on page 84
- bgp origin-as validation enable, on page 85
- bgp origin-as validation signal ibgp , on page 86
- bgp origin-as validation time , on page 87
- bgp maximum neighbor, on page 89
- bgp multipath as-path, on page 90
- bgp policy accounting, on page 91
- bgp policy propagation input flow-tag, on page 93
- bgp redistribute-internal, on page 94
- bgp router-id, on page 95
- bgp scan-time, on page 96
- bgp unsafe-ebgp-policy, on page 98
- bgp update-delay, on page 99
- bgp write-limit, on page 101
- bmp-activate, on page 103
- bmp server, on page 104
- bind-source (rpki-server), on page 107
- capability additional-paths receive, on page 108
- capability additional-paths send, on page 109
- capability orf prefix, on page 110
- capability suppress 4-byte-as, on page 112
- cef load-balancing recursive oor mode dampening-and-dlb, on page 115
- clear bgp, on page 116
- clear bgp dampening, on page 118
- clear bgp external, on page 120
- clear bgp flap-statistics, on page 122
- clear bgp nexthop performance-statistics, on page 124
- clear bgp nexthop registration, on page 126
- clear bgp peer-drops, on page 128
- clear bgp performance-statistics, on page 129
- clear bgp self-originated, on page 130
- clear bgp shutdown, on page 132
- clear bgp soft, on page 134
- clear bgp vrf rpki validation, on page 137

- cluster-id allow-equal , on page 138
- default-information originate (BGP), on page 139
- default-martian-check disable, on page 140
- default-metric (BGP), on page 141
- default-originate, on page 142
- default-prefix suppress (BGP), on page 144
- description (BGP), on page 145
- distance bgp, on page 146
- distribute bgp-ls (ISIS), on page 148
- distribute bgp-ls (OSPF), on page 149
- domain-distinguisher, on page 150
- dmz-link-bandwidth, on page 151
- dscp (BGP), on page 152
- dynamic-med interval, on page 154
- ebgp-multihop, on page 155
- export route-policy, on page 157
- export route-target, on page 158
- hw-module fib bgp-mp-pic-auto-protect-enable, on page 160
- hw-module fib dlb level-1 enable, on page 161
- hw-module profile stats j2-dynamic-stats, on page 163
- hw-module fib bgppa stats-mode, on page 164
- hw-module fib urpf disable, on page 166
- ibgp policy out enforce-modifications, on page 167
- import route-policy, on page 168
- import route-target, on page 169
- ignore-connected-check, on page 171
- graceful-maintenance, on page 172
- hw-module profile flowspec, on page 174
- hw-module fib bgppa stats-mode, on page 175
- keychain (BGP), on page 177
- keychain-disable, on page 179
- keychain inheritance-disable, on page 181
- label-allocation-mode, on page 182
- label mode, on page 184
- local-as, on page 187
- long-lived-graceful-restart, on page 189
- match flow-tag, on page 191
- maximum-paths (BGP), on page 193
- maximum-paths external-rtfilter (BGP IPv4 RT-filter address family), on page 195
- maximum-prefix (BGP), on page 196
- maximum-prefix-restart-time (BGP), on page 199
- mpls activate (BGP), on page 200
- mvpn, on page 203
- neighbor (BGP), on page 204
- neighbor-group, on page 206
- neighbor internal-vpn-client , on page 208

- network (BGP), on page 209
- network backdoor, on page 211
- next-hop-self, on page 213
- next-hop-unchanged, on page 215
- nexthop resolution prefix-length minimum, on page 217
- nexthop route-policy, on page 218
- nexthop trigger-delay, on page 220
- orf, on page 222
- password (BGP), on page 223
- password (rpki-server), on page 225
- password-disable, on page 226
- permanent-network, on page 227
- precedence, on page 228
- preference (rpki-server), on page 230
- prefix-ecmp-delay, on page 231
- purge-time (rpki-server), on page 232
- rd, on page 233
- receive-buffer-size, on page 235
- redistribute (BGP), on page 237
- refresh-time (rpki-server), on page 241
- remote-as (BGP), on page 242
- remote-bgp-label-mode-per-prefix, on page 244
- remove-private-as, on page 246
- replace as-path all, on page 248
- response-time (rpki-server), on page 250
- retain local-label, on page 251
- retain route-target, on page 252
- route-policy (BGP), on page 254
- router bgp, on page 256
- rpki server, on page 257
- selective-vrf-download disable, on page 258
- send-buffer-size, on page 259
- send-community-ebgp, on page 261
- send-extended-community-ebgp, on page 263
- session-group, on page 265
- session-open-mode, on page 267
- set flow-tag, on page 269
- set path-color external-reach, on page 270
- show bgp, on page 271
- show bgp bmp, on page 282
- show bgp update out, on page 284
- show bgp update in error process, on page 288
- show bgp update out filter-group, on page 289
- show bgp update out process, on page 290
- show bgp update out sub-group, on page 292
- show bgp update out update-group, on page 293

- show bgp vrf update in error, on page 295
- show bgp advertised, on page 296
- show bgp af-group, on page 303
- show bgp attribute-key, on page 306
- show bgp cidr-only, on page 310
- show bgp community, on page 313
- show bgp convergence, on page 318
- show bgp dampened-paths, on page 321
- show bgp flap-statistics, on page 324
- show bgp inconsistent-as, on page 329
- show bgp labels, on page 333
- show bgp large-community, on page 336
- show bgp l2vpn, on page 339
- show bgp l2vpn vpls, on page 343
- show bgp neighbor-group, on page 349
- show bgp neighbors, on page 352
- show bgp nexthops, on page 379
- show bgp paths, on page 388
- show bgp policy, on page 390
- show bgp process, on page 397
- show bgp regexp, on page 417
- show bgp route-policy, on page 421
- show bgp ipv4 rt-filter neighbor, on page 425
- show bgp session-group, on page 426
- show bgp sessions, on page 429
- show bgp summary, on page 433
- show bgp truncated-communities, on page 437
- show bgp update-group, on page 441
- show bgp vrf, on page 447
- show protocols (BGP), on page 450
- show svd role, on page 452
- show svd state, on page 453
- shutdown (BGP), on page 454
- shutdown (rpki-server), on page 456
- signalling disable, on page 457
- site-of-origin (BGP), on page 458
- slow peer (BGP router global configuration), on page 460
- slow peer (BGP neighbor address-family configuration), on page 462
- socket receive-buffer-size, on page 465
- socket send-buffer-size, on page 467
- soft-reconfiguration inbound, on page 469
- speaker-id, on page 472
- stats resource reassign , on page 473
- svd platform enable, on page 478
- table-policy, on page 480
- timers (BGP), on page 482

- timers bgp, on page 484
- transport (rpki-server), on page 486
- update limit, on page 487
- update limit address-family, on page 488
- update limit sub-group, on page 489
- update in error-handling basic disable, on page 490
- update in error-handling extended, on page 491
- update in error-handling treat-as-withdraw (BGP), on page 492
- update in labeled-unicast equivalent, on page 493
- update out logging, on page 494
- update-source, on page 495
- update wait-install, on page 497
- update wait-install delay startup, on page 498
- use, on page 499
- username (rpki-server), on page 503
- vrf (BGP), on page 504
- weight, on page 505

accept-own

To enable handling of self-originated VPN routes containing ACCEPT_OWN community attribute, use the **accept-own** command in neighbor VPNv4 or VPNv6 address family configuration mode. To disable this functionality, either use the **no** form of this command or use the command with **inheritance-disable** keyword.

accept-own [**inheritance-disable**]
no accept-own

Syntax Description

inheritance-disable	Disables handling of self-originated VPN routes containing ACCEPT_OWN community attribute and prevents inheritance of Accept Own from a parent configuration.

Command Default

Disabled

Command Modes

Neighbor address family VPNv4

Neighbor address family VPNv6

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

Task ID

Task ID	Operation
bgp	read, write

This example shows how to enable handling of accept-own community:

```
RP/0/RP0/CPU0:router#configure
RP/0/RP0/CPU0:router(config)#router bgp 100
RP/0/RP0/CPU0:router(config-bgp)#neighbor 10.2.3.4
RP/0/RP0/CPU0:router(config-bgp-nbr)#address-family vpnv4 unicast
RP/0/RP0/CPU0:router(config-bgp-nbr-af)#accept-own
```

additional-paths advertise-limit

To increase the advertisement limit of BGP additional paths from 32 to 96, use the **additional-paths advertise-limit** command in global address-family configuration mode. To restore the system to its default advertising limit (32 paths), use the **no** form of this command. To restore the default advertising limit for all neighbors belonging to a particular VRF address-family, and to prevent inheritance of advertising limit from a parent configuration, use the **disable** option.

additional-paths advertise-limit [**33-96**] [**disable**]
no additional-paths advertise-limit

Syntax Description

33-96	Configures the number of BGP additional paths. Range is 33 to 96.
disable	Restores the default advertising limit for all neighbors belonging to a specified VRF address-family, and prevents the inheritance of advertising limit from a parent configuration. **disable** keyword is applicable only to VRF address-families.

Command Default

A maximum of 32 BGP paths are supported, if the optional value from the range 33 to 96 is not specified.

Command History

Release	Modification
Release 24.2.1	This command was introduced.

Usage Guidelines

Any attempt to advertise more than the configured value will generate a warning message.

Task ID

Task ID	Operation
bgp	read, write

This example shows how to enable BGP to support advertising up to 96 paths in global address-family configuration mode.

```
Router# configure
Router(config)# router bgp 150
Router(config-bgp)# address-family ipv4 unicast
Router(config-bgp-af)# additional-paths advertise-limit 50
```

This example shows how to enable BGP to support advertising up to 96 paths in VRF address-family configuration mode.

```
Router# configure
Router(config)# router bgp 10
Router(config-bgp)# vrf foo
Router(config-bgp-vrf)# address-family ipv4 unicast
Router(config-bgp-vrf-af)# additional-paths advertise-limit 50
```

This example shows how to restore the system to its default condition using the **no additional-paths advertise-limit** command in global address-family configuration mode.

```
Router# configure
Router(config)# router bgp 150
Router(config-bgp)# address-family ipv4 unicast
Router(config-bgp-af)# no additional-paths advertise-limit
```

This example shows how to restore the system to its default condition using the **additional-paths advertise-limit disable** command in VRF address-family configuration mode.

```
Router# configure
Router(config)# router bgp 10
Router(config-bgp)# vrf foo
Router(config-bgp-vrf)# address-family ipv4 unicast
Router(config-bgp-vrf-af)# additional-paths advertise-limit disable
```

additional-paths install backup

To install a backup path into the forwarding table and provide prefix independent convergence (PIC) in case of a PE-CE link failure, use the **additional-paths install backup** command in an appropriate address family configuration mode. To prevent installing the backup path, use the **no** form of this command. To disable prefix independent convergence, use the **disable** keyword.

additional-paths install backup [**disable**]
no additional-paths install backup

Syntax Description

disable	Disables installing backup path into the forwarding table.

Command Default

None

Command Modes

VRF IPv4 address family configuration

VRF IPv6 address family configuration

VPNv4 address family configuration

VPNv6 address family configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

Task ID

Task ID	Operations
bgp	read, write

Examples

The following example shows how to enable installing a backup path into the forwarding table in VPNv4 address family mode:

```
RP/0/RP0/CPU0:router#configure
RP/0/RP0/CPU0:router(config)#router bgp 100
RP/0/RP0/CPU0:router(config-bgp)#address-family vpnv4 unicast
RP/0/RP0/CPU0:router(config-bgp-af)#additional-paths install backup
```

additional-paths receive

To configure receive capability of multiple paths for a prefix to the capable peers, use the **additional-paths receive** command in address-family configuration mode. To disable receive capability, use the **no** form of this command. To disable add-path receive capability for all neighbors belonging to a particular VRF address-family, use the **disable** option.

additional-paths receive [**disable**]
no additional-paths receive

Syntax Description

disable	Disables advertising additional paths receive capability.

Note
Use the **disable** keyword option to disable add-path receive capability for all neighbors belonging to a specified VRF address-family.

Syntax Description

This command has no keywords or arguments.

Command Default

None

Command Modes

IPv4 address family configuration

IPv6 address family configuration

VPNv4 address family configuration

VPNv6 address family configuration

VRF IPv4 address family configuration

VRF IPv6 address family configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

Use the **additional-paths receive** command to allow add-path receive capability to be negotiated for a specified address family. When the **additional-paths receive** command is configured, the receive capability is automatically enabled for all internal BGP neighbors for a specified address family. When this command is either not configured or explicitly disabled, none of the neighbors are allowed to negotiate receive capability for the address family.

After enabling the receive capability, the session needs to be reset for the configuration to take into effect.

Task ID

Task ID	Operation
bgp	read, write

This example shows how to enable additional paths receive capability under VPNv4 unicast address family:

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:routerconfig)# router bgp 100
RP/0/RP0/CPU0:router(config-bgp)# address-family vpnv4 unicast
RP/0/RP0/CPU0:router(config-bgp-af)# additional-paths receive
```

This example shows how to disable additional paths receive capability for all neighbors belonging to a particular VRF address-family (vrf1):

```
RP/0/RP0/CPU0:router#configure
RP/0/RP0/CPU0:router(config-bgp)#vrf vrf1
RP/0/RP0/CPU0:router(config-bgp-vrf)#address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-bgp-vrf-af)#additional-paths receive disable
```

additional-paths selection

To configure additional paths selection mode for a prefix, use the **additional-paths selection** command in address-family configuration mode. To disable the additional-paths selection mode for a prefix, use the **no** form of this command. To disable the additional-paths selection mode for a particular VRF address-family, use the **disable** option.

additional-paths selection {**route-policy** *route-policy-name* | **disable**}
no additional-paths selection route-policy *route-policy-name*

Syntax Description

route-policy *route-policy-name*	Specifies the name of a route policy used for additional paths selection.
disable	Disables add-path selection for a particular VRF address-family.

Command Default

None

Command Modes

IPv4 address family configuration

IPv6 address family configuration

VPNv4 address family configuration

VPNv6 address family configuration

VRF IPv4 address family configuration

VRF IPv6 address family configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

To configure additional paths selection mode for some or all prefixes, use the **additional-paths selection** command by specifying a route-policy.

Use the **additional-path selection** command with an appropriate route-policy to calculate backup paths and to enable Prefix Independent Convergence (PIC) functionality. Refer *BGP Prefix Independent Convergence Unipath Primary/Backup* section in *BGP Configuration Guide for Cisco NCS 5500 Series Routers* for details on the PIC functionality.

Task ID

Task ID	Operation
bgp	read, write

Examples

This example shows how to enable selection of additional paths:

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router bgp 100
```

```
RP/0/RP0/CPU0:router(config-bgp)# address-family vpnv4 unicast
RP/0/RP0/CPU0:router(config-bgp-af)# additional-paths selection route-policy ap1
```

This example shows how to disable add-path selection for a particular VRF address-family (vrf1):

```
RP/0/RP0/CPU0:router#configure
RP/0/RP0/CPU0:router(config-bgp)#vrf vrf1
RP/0/RP0/CPU0:router(config-bgp-vrf)#address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-bgp-vrf-af)#additional-paths selection disable
```

This example shows how to enable add-path selection for a particular VRF address-family (vrf2):

```
RP/0/RP0/CPU0:router#configure
RP/0/RP0/CPU0:router(config-bgp)#vrf vrf2
RP/0/RP0/CPU0:router(config-bgp-vrf)#address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-bgp-vrf-af)#additional-paths selection route-policy ap2
```

additional-paths send

To configure send capability of multiple paths for a prefix to the capable peers, use the **additional-paths send** command in address-family configuration mode. To disable the send capability, use the **no** form of this command.

additional-paths send [**disable**]
no additional-paths send

Syntax Description

disable Disables advertising additional paths send capability.

Note
Use the **disable** option to disable add-path send capability for all neighbors belonging to a particular VRF address-family.

Command Default

None

Command Modes

IPv4 address family configuration

IPv6 address family configuration

VPNv4 address family configuration

VPNv6 address family configuration

VRF IPv4 address family configuration

VRF IPv6 address family configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

Use the additional-paths send command to allow add-path send capability to be negotiated for a specified address family. When the **additional-paths send** command is configured, the send capability is automatically enabled for all internal BGP neighbors for the specified address family. When the command is either not configured or explicitly disabled, none of the neighbors are allowed to negotiate send capability for the address family.

After enabling the send capability, the session needs to be reset for the configuration to take into effect.

Task ID

Task ID	Operation
bgp	read, write

This example shows how to enable additional paths send capability under VPNv4 4 unicast address family:

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:routerconfig)# router bgp 100
RP/0/RP0/CPU0:router(config-bgp)# address-family vpnv4 unicast
RP/0/RP0/CPU0:router(config-bgp-af)# additional-paths send
```

This example shows how to enable add-path selection for a particular VRF address-family (vrf1):

```
RP/0/RP0/CPU0:router#configure
RP/0/RP0/CPU0:router(config-bgp)#vrf vrf1
RP/0/RP0/CPU0:router(config-bgp-vrf)#address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-bgp-vrf-af)#additional-paths send disable
```

address-family (BGP)

To enter various address family configuration modes while configuring Border Gateway Protocol (BGP), use the **address-family** command in an appropriate configuration mode. To disable support for an address family, use the **no** form of this command.

address-family {**ipv4** {**labeled-unicast** | **flowspec** | **mdt** | **multicast** | **mvpn** | **rt-filter** | **tunnel** | **unicast**} | **ipv6** {**labeled-unicast** | **multicast** | **flowspec** | **mvpn** | **unicast**} | **l2vpn vpls-vpws** | **vpnv4** {**flowspec** | **multicast** | **unicast**} | **vpnv6 {unicast | flowspec}** | **link-state link-state**}
no address-family

Syntax Description

ipv4 unicast	Specifies IP Version 4 (IPv4) unicast address prefixes.
ipv4 multicast	Specifies IPv4 multicast address prefixes.
ipv4 labeled-unicast	Specifies IPv4 labeled-unicast address prefixes. This option is available in IPv4 neighbor configuration mode and VRF neighbor configuration mode.
ipv4 tunnel	Specifies IPv4 tunnel address prefixes.
ipv4 mdt	Specifies IPv4 multicast distribution tree (MDT) address prefixes. This option is available in router configuration mode and IPv4 neighbor configuration mode.
ipv6 unicast	Specifies IP Version 6 (IPv6) unicast address prefixes.
ipv6 multicast	Specifies IP Version 6 (IPv6) multicast address prefixes.
ipv6 labeled-unicast	Specifies IPv6 labeled-unicast address prefixes. This option is available in IPv6 neighbor configuration mode.
vpnv4 unicast	Specifies VPN Version 4 (VPNv4) unicast address prefixes. This option is not available in VRF or VRF neighbor configuration mode.
vpnv6 unicast	Specifies VPN Version 6 (VPNv6) unicast address prefixes. This option is not available in VRF or VRF neighbor configuration mode.
l2vpn vpls-vpws	Specifies L2VPN vpls-vpws address prefixes.
ipv4 rt-filter	Specifies IPv4 rt-filter address prefixes.
ipv4 mvpn	Specifies IPv4 mvpn address prefixes.
ipv6 mvpn	Specifies IPv6 mvpn address prefixes.
link-state link-state	Advertises link-state database of a network via BGP.
flowspec	Specifies flowspec configuration mode.
vpnv4 multicast	Specifies VPNv4 multicast prefixes.

Command Default

An address family must be explicitly configured in the router configuration mode for the address family to be active in BGP. Similarly, an address family must be configured under the neighbor for the BGP session to

be established for that address family. An address family must be configured in router configuration mode before it can be configured under a neighbor.

Command Modes

Router configuration

Neighbor configuration

Neighbor group configuration

Flowspec configuration

VRF configuration

VRF neighbor configuration (IPv4 address families)

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

Use the **address-family** command to enter various address family configuration modes while configuring BGP routing sessions. When you enter the **address-family** command from router configuration mode, you enable the address family and enter global address family configuration mode.

The IPv4 unicast address family must be configured in router configuration mode before configuring the IPv4 labeled-unicast address family for a neighbor in neighbor configuration mode.

The IPv4 unicast address family must be configured in router configuration mode before configuring the IPv4 labeled-unicast address family for a neighbor in neighbor configuration mode. The IPv6 unicast address family must be configured in router configuration mode before configuring the IPv6 labeled-unicast address family for a neighbor in neighbor configuration mode.

Table 1: Address Family Submode Support

Address Family	Supported in Router Submode	Supported in Neighbor Submode	Comments
ipv4 unicast	yes	yes	—
ipv4 multicast	yes	yes	—
ipv4 mdt	yes	yes	—
ipv4 tunnel	yes	yes	—
ipv4 labeled-unicast	no	yes	The ipv4 labeled-unicast address family can be configured only as a neighbor address family; however, it requires that the ipv4 unicast address family be configured as the router address family first.
vpnv4 unicast	yes	yes	—
ipv6 unicast	yes	yes	—
ipv6 multicast	yes	yes	—

Address Family	Supported in Router Submode	Supported in Neighbor Submode	Comments
ipv6 labeled-unicast	no	yes	The ipv6 labeled-unicast address family can be configured only as a neighbor address family; however, it requires that the ipv6 unicast address family be configured as the router address family first.
vpnv6 unicast	yes	yes	—
l2vpn vpls-vpws	yes	yes	—
ipv4 rt-filter	yes	yes	—
ipv4 mvpn	yes	yes	—
ipv6 mvpn	yes	yes	—
link-state	yes	yes	—
flowspec	yes	yes	If you configure flowspec redirect next hop rule, IPv4 and IPv6 packets drop when there is an invalid next hop.

When you enter the **address-family** command from neighbor configuration mode, you activate the address family on the neighbor and enter neighbor address family configuration mode. IPv4 neighbor sessions support IPv4 unicast, multicast , and labeled-unicast, and VPNv4 unicast address families. IPv6 neighbor sessions support IPv6 unicast and multicast address families.

Outbound Route Filter (ORF) capability is not supported with address-family l2vpn vpls-vpws

Task ID

Task ID	Operations
bgp	read, write

Examples

The following example shows how to place the router in global address family configuration mode for the IPv4 address family:

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router bgp 100
RP/0/RP0/CPU0:router(config-bgp)# address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-bgp-af)#
```

The following example shows how to activate IPv4 multicast for neighbor 10.0.0.1 and place the router in neighbor address family configuration mode for the IPv4 multicast address family:

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router# router bgp 1
RP/0/RP0/CPU0:router(config-bgp)# address-family ipv4 multicast
RP/0/RP0/CPU0:router(config-bgp-af)# exit
```

```
RP/0/RP0/CPU0:router(config-bgp)# neighbor 10.0.0.1
RP/0/RP0/CPU0:router(config-bgp-nbr)# remote-as 1
RP/0/RP0/CPU0:router(config-bgp-nbr)# address-family ipv4 multicast
RP/0/RP0/CPU0:router(config-bgp-nbr-af)#
```

The following example shows how to place the router in global address family configuration mode for the IPv4 tunnel address family:

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router bgp 12
RP/0/RP0/CPU0:router(config-bgp)# address-family ipv4 tunnel
RP/0/RP0/CPU0:router(config-bgp-af)#
```

The following example shows how to place the router in global address family link-state configuration mode:

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router bgp 100
RP/0/RP0/CPU0:router(config-bgp)# address-family link-state link-state
RP/0/RP0/CPU0:router(config-bgp-af)#
```

The following example shows how to exchange link-state information with a BGP neighbor:

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router bgp 100
RP/0/RP0/CPU0:router(config-bgp)# neighbor 10.0.0.2
RP/0/RP0/CPU0:router(config-bgp-nbr)# remote-as 1
RP/0/RP0/CPU0:router(config-bgp-nbr)# address-family link-state link-state
RP/0/RP0/CPU0:router(config-bgp-nbr-af)#
```

The following example shows how to place the router in flowspec sub-address family configuration mode for the IPv4 address family:

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router bgp 100
RP/0/RP0/CPU0:router(config-bgp)# address-family ipv4 flowspec
RP/0/RP0/CPU0:router(config-bgp-af)#
```

advertise best-external

To advertise the best–external path to the iBGP and route-reflector peers, when a locally selected bestpath is from an internal peer, use the **advertise best-external** command in an appropriate address family configuration mode. To prevent advertising the best–external path, use the **no** form of this command. To disable advertising the best–external path, use the **disable** keyword.

advertise best-external [**disable**]
no advertise best-external

Syntax Description

disable	Disables best–external configuration for the VRF.

Command Default

None

Command Modes

VRF IPv4 address family configuration

VRF IPv6 address family configuration

L2VPN address family configuration

VPNv4 address family configuration

VPNv6 address family configuration

IPv4 address family configuration

IPv6 address family configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

Task ID

Task ID	Operations
bgp	read, write

Examples

The following example shows how to enable advertising the best–external path VPNv4 unicast address family mode:

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router bgp 100
RP/0/RP0/CPU0:router(config-bgp)# address-family vpnv4 unicast
RP/0/RP0/CPU0:router(config-bgp-af)# advertise best-external
```

advertise permanent-network

To identify the peers to whom the permanent paths must be advertised, use the **advertise permanent-network** command in the neighbor address family configuration mode. To stop advertising the permanent p, use the **no** form of this command. The permanent paths will always be advertised to peers having advertise permanent-network configuration, even if a different best-path is available. The permanent path is not advertised to peers that are not configured to receive permanent path.

The permanent path supports only prefixes in IPv4 unicast and IPv6 unicast address-families under the default Virtual Routing and Forwarding (VRF).

advertise permanent-network
no advertise permanent-network

Syntax Description

This command has no arguments or keywords.

Command Modes

Neighbor address-family configuration.

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

Task ID

Task ID	Operations
bgp	read, write

Examples

This example shows how to advertise permanent path:

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router bgp 100
RP/0/RP0/CPU0:router(config-bgp)# neighbor 10.1.1.1
RP/0/RP0/CPU0:router(config-bgp-nbr)# remote-as 4713
RP/0/RP0/CPU0:router(config-bgp-nbr)# address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-bgp-nbr-af)# advertise permanent-network
```

advertisement-interval

To set the minimum interval between the sending of Border Gateway Protocol (BGP) routing updates, use the **advertisement-interval** command in an appropriate configuration mode. To remove the **advertisement-interval** command from the configuration file and restore the system to its default interval values, use the **no** form of this command.

advertisement-interval *seconds*
no advertisement-interval [*seconds*]

Syntax Description

seconds	Minimum interval between sending BGP routing updates (in seconds). Range is 0 to 600.

Command Default

Default minimum interval:

For internal BGP (iBGP) peers is 0 seconds

For external BGP (eBGP) peers is 30 seconds

For customer edge (CE) peers is 0 seconds

Command Modes

Neighbor configuration

Neighbor group configuration

Session group configuration

VRF neighbor configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

If this command configures a neighbor group or session group, all neighbors using the group inherit the configuration. Values of commands configured specifically for a neighbor override inherited values.

Task ID

Task ID	Operations
bgp	read, write

Examples

The following example shows how to set the minimum time between sending BGP routing updates to 10 seconds:

```
RP/0/RP0/CPU0:router(config)# router bgp 5
RP/0/RP0/CPU0:router(config-bgp)# neighbor 10.1.1.1
RP/0/RP0/CPU0:router(config-bgp-nbr)# remote-as 100
RP/0/RP0/CPU0:router(config-bgp-nbr)# advertisement-interval 10
```

advertise vpn-imported label-mode per-vrf

To enable per-VRF label allocation for imported VPN routes, use the **advertise vpn-imported label-mode per-vrf** command in VRF address-family configuration mode.

advertise vpn-imported label-mode per-vrf

Command Default

Per-prefix is the default label allocation mode.

Command Modes

VRF address-family configuration

Command History

Release	Modification
Release 24.4.1	This command was introduced.

Usage Guidelines

Use this command to override the default per-prefix label allocation behavior and enforce per-VRF label allocation.

Task ID

Task ID	Operation
bgp	read, write

Examples

This example shows how to enable per-VRF label allocation.

```
Router# configure
Router(config)# vrf vrf_1
Router(config-vrf)# address-family ipv4 unicast
Router(config-vrf-af)# advertise vpn-imported label-mode per-vrf
```

af-group

To create an address family group for Border Gateway Protocol (BGP) neighbors and enter address family group configuration mode, use the **af-group** command in XR Config mode. To remove an address family group, use the **no** form of this command.

af-group *af-group-name* **address-family**
no af-group

Syntax Description

af-group-name	Address family group name.
address-family	Enters address family configuration mode.
ipv4 unicast	Specifies IP Version 4 (IPv4) unicast address prefixes.
ipv4 multicast	Specifies IPv4 multicast address prefixes.
ipv4 labeled-unicast	Specifies IPv4 labeled unicast address prefixes.
ipv4 tunnel	Specifies IPv4 tunnel address prefixes.
ipv4 mdt	Specifies IPv4 multicast distribution tree (MDT) address prefixes.
ipv6 unicast	Specifies IP Version 6 (IPv6) unicast address prefixes.
ipv6 multicast	Specifies IPv6 multicast address prefixes.
ipv6 labeled-unicast	Specifies IPv6 labeled unicast address prefixes.
vpnv4 unicast	Specifies VPN Version 4 (VPNv4) unicast address prefixes.
vpnv6 unicast	Specifies VPN Version 6 (VPNv6) unicast address prefixes.

Command Default

No BGP address family group is configured.

Command Modes

XR Config mode

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

Use the **af-group** command to group address family-specific neighbor commands within an IPv4 or IPv6 address family. Neighbors that have address family configuration are able to use the address family group. Further, neighbors inherit the configuration parameters of the entire address family group.

You cannot define two address family groups with the same name in different address families.

Note A non-existent af-group maybe specified under a neighbor, but will only take effect when the af-group is configured.

Task ID

Task ID	Operations
bgp	read, write

Examples

The following example shows how to create address family group group1 and enter address family group configuration mode for IPv4 unicast. Group1 contains the next-hop-self feature, which is inherited by neighbors that use address family group1.

```
RP/0/RP0/CPU0:router(config)# router bgp 100
RP/0/RP0/CPU0:router(config-bgp)# af-group group1 address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-bgp-afgrp)# next-hop-self
```

aggregate-address

To create an aggregate entry in a Border Gateway Protocol (BGP) routing table, use the **aggregate-address** command in an appropriate configuration mode. To remove the **aggregate-address** command from the configuration file and restore the system to its default condition, use the **no** form of this command.

aggregate-address *address/mask-length* **[as-set] [as-confed-set] [summary-only] [route-policy** *route-policy-name***]**
no aggregate-address

Syntax Description

address	Aggregate address.
/mask-length	Aggregate address mask length.
as-set	(Optional) Generates autonomous system set path information and community information from contributing paths.
as-confed-set	(Optional) Generates autonomous system confederation set path information from contributing paths.
summary-only	(Optional) Filters all more-specific routes from updates.
route-policy *route-policy-name*	(Optional) Specifies the name of a route policy used to set the attributes of the aggregate route.

Command Default

When you do not specify this command, no aggregate entry is created in the BGP routing table.

Command Modes

IPv4 address family configuration

IPv6 address family configuration

VRF IPv4 address family configuration

VRF IPv6 address family configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

You can implement aggregate routing in BGP either by redistributing an aggregate route into BGP using the **network** command or the **aggregate-address** command.

Use the **aggregate-address** command without optional arguments to create an aggregate entry in the BGP routing table if any more-specific BGP routes are available that fall in the specified range. The aggregate route is advertised as coming from your autonomous system and has the atomic aggregate attribute set to show that information might be missing. (By default, the atomic aggregate attribute is set unless you specify the **as-set** keyword.)

Use of the **as-set** keyword creates an aggregate entry using the same rules that the command follows without this keyword. However, the advertised path for this route is an AS_SET, a set of all autonomous systems contained in all paths that are being summarized.

Do not use this form of the **aggregate-address** command when aggregating many paths because this route must be continually withdrawn and updated as autonomous system path reachability information for the summarized routes changes.

Use the **as-confed-set** keyword to create an AS_CONFED_SET in the autonomous system path of the aggregate from any confederation segments in the paths being summarized. This keyword takes effect only if the **as-set** keyword is also specified.

Use of the **summary-only** keyword creates an aggregate entry (for example, 10.0.0.0/8) but suppresses advertisements of more-specific routes to all neighbors. If you want to suppress only advertisements to certain neighbors, use the **route-policy (BGP)** command in neighbor address family configuration mode with caution. If a more-specific route leaks out, all BGP speakers (the local router) prefer that route over the less-specific aggregate you generate (using longest-match routing).

Use the **route-policy** keyword to specify a routing policy for the aggregate entry. The **route-policy** keyword is used to select which more-specific information to base the aggregate entry on and which more-specific information to suppress. You can also use the keyword to modify the attributes of the aggregate entry.

Task ID

Task ID	Operations
bgp	read, write

Examples

The following example shows how to create an aggregate address. The path advertised for this route is an autonomous system set consisting of all elements contained in all paths that are being summarized.

```
RP/0/RP0/CPU0:router(config)# router bgp 100
RP/0/RP0/CPU0:router(config-bgp)# address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-bgp-af)# aggregate-address 10.0.0.0/8 as-set
```

aigp

To enable sending and receiving of accumulated interior gateway protocol (AiGP) attribute per eBGP neighbor, use the **aigp**command in appropriate configuration mode. To disable this functionality, either use the **disable** keyword or use the **no** form of this command.

aigp [**disable**]
no aigp

Syntax Description

disable	Disables sending or receiving AiGP attribute.

Command Default

Send or recive of AiGP attribute is disabled for eBGP neighbors

Command Modes

IPv4 address family configuration

IPv6 address family configuration

VRF IPv4 address family configuration

VRF IPv6 address family configuration

VPNv4 address family configuration

VPNv6 address family configuration

Neighbor address family configuration

VRF neighbor address family configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

Task ID

Task ID	Operation
bgp	read, write

Examples

The following example shows how to enable AiGP send and receive capability under neighbor address family (IPv4 unicast):

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router bgp 100
RP/0/RP0/CPU0:router(config-bgp)# neighbor 10.2.3.4
RP/0/RP0/CPU0:router(config-bgp-nbr)# address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-bgp-nbr-af)# aigp
```

aigp send-cost-community

To send Accumulated Interior Gateway Protocol (AiGP) value in cost community, use the **aigp send-cost-community** command in appropriate configuration mode. To disable sending AiGP value in cost community, either use the **no** form of this command or the **disable** keyword.

aigp send-cost-community {*cost-id* | **disable**} **poi** {**igp-cost** | **pre-bestpath**} [**transitive**]
no aigp send-cost-community

Syntax Description

cost-comm-id	Specifies the Cost community ID. The range is 0 to 255.
poi	Point of insertion for bestpath calculation.
igp-cost	Configures that cost community be used after iGP distance to next hop.
pre-bestpath	Configures cost community as first step in best path calculation.
transitive	(Optional) Enables transitive cost community
disable	Disables sending AiGP value in cost community.

Command Default

Sending AiGP value in cost community is disabled

Command Modes

Neighbor address family configuration

VRF neighbor address family configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

Cost community point of insertion can be configured either to be pre-bestpath or after igp cost. The **transitive** keyword is not required for iBGP sessions. However, the **transitive** keyword is required for eBGP sessions to convert AiGP metric into cost-community and advertise to the eBGP neighbors.

Task ID

Task ID	Operation
bgp	read, write

Examples

The following example shows how to enable sending AiGP value in cost community ID 254 under neighbor address family (IPv4 unicast):

```
RP/0/RP0/CPU0:router# configure
```

```
RP/0/RP0/CPU0:router(config)# router bgp 100
RP/0/RP0/CPU0:router(config-bgp)# neighbor 10.2.3.4
RP/0/RP0/CPU0:router(config-bgp-nbr)# address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-bgp-nbr-af)# aigp send-cost-community 254
```

allocate-label

To allocate Multiprotocol Label Switching (MPLS) labels for specific IPv4 unicast or IPv6 unicast or VPN routing and forwarding (VRF) IPv4 unicast routes so that the BGP router can send labels with BGP routes to a neighboring router configured for labeled- or VPN routing and forwarding (VRF) IPv6 unicast sessions, use the **allocate-label** command in the appropriate configuration mode. To restore the system to its default condition, use the **no** form of this command.

allocate-label {**route-policy** *route-policy-name* | **all**}
no allocate-label {**route-policy** *route-policy-name* | **all**}

Syntax Description

all	Allocates labels for all prefixes
route-policy *route-policy-name*	Uses a route policy to select prefixes for label allocation.

Command Default

No default behavior or values

Command Modes

IPv4 address family configuration

IPv6 address family configuration

VRF IPv4 address family configuration

VRF IPv6 address family configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

Use the **allocate-label** command with a route policy to trigger BGP to allocate labels for all or a filtered set of global IPv4 routes (as dictated by the route policy). The command enables autonomous system border routers (ASBRs) that have labeled IPv4 unicast sessions to exchange Multiprotocol Label Switching (MPLS) labels with the IPv4 routes to the other autonomous system (AS) in Layer 3 Virtual Private Network (L3VPN) inter-AS deployments.



Note The **allocate-label all** command is functionally equivalent to the **allocate-label route-policy** *route-policy-name* command when the route policy is a pass-all policy.

See the *L3VPN Configuration Guide for Cisco NCS 5500 Series Routers* for information on using the **allocate-label** command for L3VPN inter-AS deployments and carrier-supporting-carrier IPv4 BGP label distribution.

Task ID

Task ID	Operations
bgp	read, write

Examples

The following example shows how to enable allocating labels for IPv4 routes:

```
RP/0/RP0/CPU0:router(config)# router bgp 6
RP/0/RP0/CPU0:router(config-bgp)# address family ipv4 unicast
RP/0/RP0/CPU0:router(config-bgp-af)# allocate-label route-policy policy_A
```

allowas-in

To allow an AS path with the provider edge (PE) autonomous system number (ASN) a specified number of times, use the **allowas-in** command in an appropriate configuration mode. To restore the system to its default condition, use the **no** form of this command.

allowas-in [*as-occurrence-number*]
no allowas-in [*as-occurrence-number*]

Syntax Description

as-occurrence-number	(Optional) Number of times a PE ASN is allowed. Range is 1 to 10.

Command Default

No default behavior or values

Command Modes

Address family group configuration

Neighbor address family configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

Hub and spoke VPN networks require looping back of routing information to the hub PE through the hub customer edge (CE). See *MPLS Configuration Guide for Cisco NCS 5500 Series Routers* for information on hub and spoke VPN networks. This looping back, in addition to the presence of the PE ASN, causes the looped-back information to be dropped by the hub PE.

The **allowas-in** command prevents the looped-back information from being dropped by replacing the neighbor autonomous system number (ASN) with the PE ASN in the AS path. This allows the VPN customer to see a specified number of occurrences of the PE ASN in the AS path.

Task ID

Task ID	Operations
bgp	read, write

Examples

The following example shows how to allow five occurrences of the PE ASN:

```
RP/0/RP0/CPU0:router(config)# router bgp 6
RP/0/RP0/CPU0:router(config-bgp)# af-group group_1 address-family vpnv4 unicast
RP/0/RP0/CPU0:router(config-bgp-afgrp)# allowas-in 5
```

allowconfedas-in

To enable peer routers in the same confederation to learn from each other for a specified number of times, use the **allowconfedas-in** command in neighbor address family configuration mode.

allowconfedas-in *confederation-as-count*

Syntax Description

confederation-as-count	Number of times a provider edge (PE) confederation AS number is allowed. Range: 1 to 10. Default: 3

Command Default

Peers within the same Confederation will not learn each other's routes.

Command Modes

Neighbor address-family configuration

Command History

Release	Modification
Release 7.11.1	This command was introduced.

Task ID

Task ID	Operation
bgp	read, write

Examples

In this example, the peer routers in the same confederation learn from each other thrice:

```
Router# router bgp 65001
Router(config-bgp)# bgp confederation peers 65002
Router(config-bgp)# bgp confederation identifier 100
Router(config-bgp)# neighbor 198.51.100.3
Router(config-bgp-nbr)# address-family ipv4 unicast
Router(config-bgp-nbr-af)# allowconfedas-in 1
```

as-format

To configure the router's Autonomous system number (ASN) notation to asdot format, use the as-format command in XR Config mode. To restore the system to its default condition, use the **no** form of this command.

as-format asdot
no

Syntax Description

asdot	Specifies the Autonomous system number (ASN) notation to asdot format.

Command Default

The default value, if the as-format command is not configured, is asplain.

Command Modes

XR Config mode

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

Task ID

Task ID	Operations
bgp	read, write

Examples

The following example shows how to configure the ASN notation to the asdot format:

```
RP/0/RP0/CPU0:router(config)# as-format asdot
```

as-override

To configure a provider edge (PE) router to override the autonomous system number (ASN) of a site with the ASN of a provider, use the **as-override** command which works for both VRF and non-VRF neighbor address family configuration mode. To restore the system to its default condition, use the **no** form of this command.

as-override [**inheritance-disable**]
no as-override [**inheritance-disable**]

Syntax Description

inheritance-disable	(Optional) Prevents the **as-override** command from being inherited from a parent group.

Command Default

Automatic override of the ASN is disabled.

Command Modes

VRF and non-VRF neighbor address family configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

Use the **as-override** command in conjunction with the site-of-origin (SoO) feature, identifying the site where a route originated, and preventing routing loops between routers within a VPN.

Task ID

Task ID	Operations
bgp	read, write

Examples

The following example shows how to configure an ASN override:

```
RP/0/RP0/CPU0:router(config)# router bgp 6
RP/0/RP0/CPU0:router(config-bgp)# vrf vrf_A
RP/0/RP0/CPU0:router(config-bgp-vrf)# neighbor 192.168.70.24
RP/0/RP0/CPU0:router(config-bgp-vrf-nbr)# remote-as 10
RP/0/RP0/CPU0:router(config-bgp-vrf-nbr)# address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-bgp-vrf-nbr-af)# as-override
```

as-path-loopcheck out disable

To disable AS PATH loop checking for outbound updates, use the **as-path-loopcheck out disable** command in an appropriate address family configuration mode. To re-enable the default AS PATH loop checking, use the **no** form of this command.

as-path-loopcheck out disable
no as-path-loopcheck out disable

Syntax Description

This command has no keywords or arguments.

Command Default

AS PATH loop checking for outbound updates is enabled if there is only one neighbor and disabled if there are multiple neighbors in the update group.

Command Modes

IPv4 address family

IPv6 address family

L2VPN address family

VPNv4 address family

VPNv6 address family

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

Configure the **as-path-loopcheck out disable** command to disable the default behavior of PE router not announcing BGP routes to the CE router if the routes contain an AS number matching the AS number of the receiving CE router.

Task ID

Task ID	Operation
bgp	read, write

This example shows how to configure **as-path-loopcheck out disable** under IPv6 unicast address family:

```
RP/0/RP0/CPU0:router#configure
RP/0/RP0/CPU0:router(config)#router bgp 100
RP/0/RP0/CPU0:router(config-bgp)#address-family ipv6 unicast
RP/0/RP0/CPU0:router(config-bgp-af)#as-path-loopcheck out disable
```

attribute large-community

To specify the large community BGP path attributes in the created attribute filter and the action to be taken on the update message received from the BGP neighbor, use the **attribute large-community** command in the BGP attribute-filter group command mode.

attribute large-community { **discard** | **treat-as-withdraw** }

Syntax Description

discard	Specifies that the large-community BGP path attribute received in a BGP update message from the neighbor must be discarded.
treat-as-withdraw	Indicates that the entire BGP update message that contains the large-community attribute has to be treated as a 'withdraw' message for the Network Layer Reachability Information (NLRI) contained in the update message.

Command Default

None

Command Modes

BGP attribute-filter group mode

Command History

Release	Modification
Release 6.3.1	This command was introduced.

Usage Guidelines

None.

Task ID

Task ID	Operation
bgp	read, write

The following example creates the attribute-filter named dogbert and specifies the large-community attribute and the action to discard. The dogbert attribute-filter is applied to the BGP neighbor 10.10.10.3. If the large-community BGP path attribute is received in a BGP update message from the neighbor 10.10.10.3, then the attribute will be discarded before further processing of the message.

```
RP/0/RP0/CPU0:router#configure
RP/0/RP0/CPU0:router(config)#router bgp 3
RP/0/RP0/CPU0:router(config-bgp)#attribute-filter group dogbert
RP/0/RP0/CPU0:router(config-bgp-attrfg)#attribute large-community discard
RP/0/RP0/CPU0:router(config-bgp-attrfg)#neighbor 10.10.10.3
RP/0/RP0/CPU0:router(config-bgp-nbr)#update in filtering
RP/0/RP0/CPU0:router(config-nbr-upd-filter)#attribute-filter group dogbert
```

attribute-filter group

To configure attribute-filter group command mode, use the attribute-filter group command in an appropriate configuration mode. To disable attribute-filter group command mode, use the no form of this command.

attribute-filter group *group-name*
no attribute-filter group *group-name*

Syntax Description

group-name	Specifies the name of the attribute-filter group.

Command Default

Attribute-filter group command mode is disabled.

Command Modes

Router configuration

Neighbor configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

Use the **attribute-filter group** command in neighbor configuration mode to configure a specific attribute filter group for a BGP neighbor.

Task ID

Task ID	Operation
bgp	read, write

This example shows how to configure the attribute-filter group command mode:

```
RP/0/RP0/CPU0:router#configure
RP/0/RP0/CPU0:router(config)#router bgp 100
RP/0/RP0/CPU0:router(config-bgp)#attribute-filter group ag_discard_med
RP/0/RP0/CPU0:router(config-bgp-attrfg)#
```

This example shows how to configure the attribute filter group for a BGP neighbor:

```
RP/0/RP0/CPU0:router#configure
RP/0/RP0/CPU0:router(config)#router bgp 100
RP/0/RP0/CPU0:router(config-bgp)#neighbor 10.0.1.101
RP/0/RP0/CPU0:router(config-bgp-nbr)#remote-as 6461
RP/0/RP0/CPU0:router(config-bgp-nbr)#update in filtering
RP/0/RP0/CPU0:router(config-nbr-upd-filter)#attribute-filter group ag_discard_med
```

bfd (BGP)

To specify a bidirectional forwarding detection (BFD) **multiplier** and **minimum-interval** arguments per neighbor, use the **bfd** command in neighbor address family independent configuration mode. To return to the system defaults, use the **no** form of this command.

Previous to this enhancement, BFD could be configured only in global scope in BGP. This change makes available two new command-line arguments under neighbor address family independent configuration:

bfd {**multiplier** | **minimum-interval**} *value*
no bfd {**multiplier** | **minimum-interval**} *value*

Syntax Description

multiplier	*value*	Specifies the BFD session's multiplier value for the neighbor.
minimum-interval	*value*	Specifies the BFD session's minimum-interval value for the neighbor.

Command Default

No default per neighbor parameters are set.

Command Modes

Neighbor address family independent configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

If the minimum interval is changed using the **bfd minimum-interval**command, the new parameter updates all affected BFD sessions under the command mode in which the minimum interval was changed.

If the multiplier is changed using the **bfd multiplier** command, the new parameter is used to update only the BFD sessions associated with the affected neighbor gets affected.

The assumption is that when BFD fast-detect is enabled under neighbor address family independent configuration, the values for the **multiplier** and **minimum-interval** values are always derived from the per-neighbor values if they are configured; otherwise, they are to be taken from the global BGP configuration mode. In the event that this has not been explicitly stated, then these values are taken to be the default values. Also, the **bfd** arguments can be configured under neighbor-group and session-group and the inheritance adheres to the standard way of BGP configuration inheritance.

Accordingly, there are four cases in which bfd-fast detect is enabled.

This is shown in table below where the BFD value is either multiplier or minimum-interval. Local indicates per NBR value, global is the BGP global value.

BFD value (global)	BFD value (local)	Result
Yes	Yes	BFD value (local)
Yes	No	BFD value (global)
No	Yes	BFD value (local)
No	No	BFD value (default)

Examples

The following example shows how to specify the BFD session's multiplier value for the neighbor:

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router bgp 65000
RP/0/RP0/CPU0:router(config-bgp-nbrgrp)#neighbor 3.3.3.2
RP/0/RP0/CPU0:router(config-bgp-nbr)# bfd minimum-interval 311
RP/0/RP0/CPU0:router(config-bgp-nbr)# bfd multiplier 7
RP/0/RP0/CPU0:router(config-bgp-nbr)# neighbor 5.5.5.2
RP/0/RP0/CPU0:router(config-bgp-nbr)# bfd minimum-interval 318
RP/0/RP0/CPU0:router(config-bgp-nbr)# bfd multiplier 4
RP/0/RP0/CPU0:router(config-bgp-nbr)# vrf one
RP/0/RP0/CPU0:router(config-bgp-vrf)# neighbor 3.12.1.2
RP/0/RP0/CPU0:router(config-bgp-vrf-nbr)# bfd minimum-interval 119
RP/0/RP0/CPU0:router(config-bgp-vrf-nbr)# bfd multiplier 10
RP/0/RP0/CPU0:router(config-bgp-vrf-nbr)# commit
```

```
RP/0/RP0/CPU0:router# show bfd session
Interface            Dest Addr           Local det time(int*mult)      State
                                           Echo            Async
-------------------- --------------- ---------------- ---------------- ---------
Gi0/2/0/2            3.3.3.2         2177ms(311ms*7)  14s(2s*7)        UP
Gi0/2/0/2.1          3.12.1.2        1190ms(119ms*10) 20s(2s*10)       UP
PO0/3/0/6            5.5.5.2         1272ms(318ms*4)  8s(2s*4)         UP


RP/0/RP0/CPU0:router# show bfd session detail
I/f: GigabitEthernet0/2/0/2, Location: 0/2/CPU0, dest: 3.3.3.2, src: 3.3.3.1
 State: UP for 0d:0h:4m:44s, number of times UP: 1
Received parameters:
 Version: 1, desired tx interval: 2 s, required rx interval: 2 s
 Required echo rx interval: 1 ms, multiplier: 7, diag: None
 My discr: 524295, your discr: 524296, state UP, D/F/P/C/A: 0/0/0/1/0
Transmitted parameters:
 Version: 1, desired tx interval: 2 s, required rx interval: 2 s
 Required echo rx interval: 1 ms, multiplier: 7, diag: None
 My discr: 524296, your discr: 524295, state UP, D/F/P/C/A: 0/0/0/1/0
Timer Values:
 Local negotiated async tx interval: 2 s
 Remote negotiated async tx interval: 2 s
 Desired echo tx interval: 311 ms, local negotiated echo tx interval: 311 ms
 Echo detection time: 2177 ms(311 ms*7), async detection time: 14 s(2 s*7)
Local Stats:
 Intervals between async packets:
   Tx: Number of intervals=100, min=1664 ms, max=2001 ms, avg=1838 ms
       Last packet transmitted 313 ms ago
   Rx: Number of intervals=100, min=1662 ms, max=2 s, avg=1828 ms
       Last packet received 1615 ms ago
 Intervals between echo packets:
   Tx: Number of intervals=100, min=181 ms, max=462 ms, avg=229 ms
       Last packet transmitted 289 ms ago
   Rx: Number of intervals=100, min=178 ms, max=461 ms, avg=229 ms
       Last packet received 287 ms ago
 Latency of echo packets (time between tx and rx):
   Number of packets: 100, min=0 us, max=4 ms, avg=860 us
Session owner information:
  Client           Desired interval         Multiplier
  ---------------- --------------------     --------------
  bgp-0            311 ms                   7

I/f: GigabitEthernet0/2/0/2.1, Location: 0/2/CPU0, dest: 3.12.1.2, src: 3.12.1.1
 State: UP for 0d:0h:4m:44s, number of times UP: 1
Received parameters:
```

```
 Version: 1, desired tx interval: 2 s, required rx interval: 2 s
 Required echo rx interval: 1 ms, multiplier: 10, diag: None
 My discr: 524296, your discr: 524295, state UP, D/F/P/C/A: 0/0/0/1/0
Transmitted parameters:
 Version: 1, desired tx interval: 2 s, required rx interval: 2 s
 Required echo rx interval: 1 ms, multiplier: 10, diag: None
 My discr: 524295, your discr: 524296, state UP, D/F/P/C/A: 0/0/0/1/0
Timer Values:
 Local negotiated async tx interval: 2 s
 Remote negotiated async tx interval: 2 s
 Desired echo tx interval: 119 ms, local negotiated echo tx interval: 119 ms
 Echo detection time: 1190 ms(119 ms*10), async detection time: 20 s(2 s*10)
Local Stats:
 Intervals between async packets:
   Tx: Number of intervals=100, min=1664 ms, max=2001 ms, avg=1838 ms
       Last packet transmitted 314 ms ago
   Rx: Number of intervals=100, min=1662 ms, max=2 s, avg=1828 ms
       Last packet received 1616 ms ago
 Intervals between echo packets:
   Tx: Number of intervals=100, min=120 ms, max=223 ms, avg=125 ms
       Last packet transmitted 112 ms ago
   Rx: Number of intervals=100, min=119 ms, max=223 ms, avg=125 ms
       Last packet received 110 ms ago
 Latency of echo packets (time between tx and rx):
   Number of packets: 100, min=0 us, max=2 ms, avg=850 us
Session owner information:
  Client           Desired interval         Multiplier
  ---------------- --------------------     --------------
  bgp-0            119 ms                   10

I/f: GigabitEthernet0/3/0/6, Location: 0/3/CPU0, dest: 5.5.5.2, src: 5.5.5.1
 State: UP for 0d:0h:4m:50s, number of times UP: 1
Received parameters:
 Version: 1, desired tx interval: 2 s, required rx interval: 2 s
 Required echo rx interval: 1 ms, multiplier: 4, diag: None
 My discr: 786436, your discr: 786433, state UP, D/F/P/C/A: 0/0/0/1/0
Transmitted parameters:
 Version: 1, desired tx interval: 2 s, required rx interval: 2 s
 Required echo rx interval: 1 ms, multiplier: 4, diag: None
 My discr: 786433, your discr: 786436, state UP, D/F/P/C/A: 0/0/0/1/0
Timer Values:
 Local negotiated async tx interval: 2 s
 Remote negotiated async tx interval: 2 s
 Desired echo tx interval: 318 ms, local negotiated echo tx interval: 318 ms
 Echo detection time: 1272 ms(318 ms*4), async detection time: 8 s(2 s*4)
Local Stats:
 Intervals between async packets:
   Tx: Number of intervals=100, min=1663 ms, max=2 s, avg=1821 ms
       Last packet transmitted 1740 ms ago
   Rx: Number of intervals=100, min=1663 ms, max=2001 ms, avg=1832 ms
       Last packet received 160 ms ago
 Intervals between echo packets:
   Tx: Number of intervals=100, min=181 ms, max=484 ms, avg=232 ms
       Last packet transmitted 44 ms ago
   Rx: Number of intervals=100, min=179 ms, max=484 ms, avg=232 ms
       Last packet received 41 ms ago
 Latency of echo packets (time between tx and rx):
   Number of packets: 100, min=0 us, max=3 ms, avg=540 us
Session owner information:
  Client           Desired interval         Multiplier
  ---------------- --------------------     --------------
  bgp-0            318 ms                   4

RP/0/RP0/CPU0:router# show bgp nei 3.3.3.2
```

```
 BGP neighbor is 3.3.3.2
  Remote AS 500, local AS 65000, external link
  Remote router ID 16.0.0.1
   BGP state = Established, up for 00:05:01
   BFD enabled (session up): mininterval: 311 multiplier: 7
   Last read 00:00:56, hold time is 180, keepalive interval is 60 seconds
   Precedence: internet
   Neighbor capabilities:
     Route refresh: advertised and received
     4-byte AS: advertised and received
     Address family IPv4 Unicast: advertised and received
   Received 8 messages, 0 notifications, 0 in queue
   Sent 9 messages, 1 notifications, 0 in queue
   Minimum time between advertisement runs is 30 seconds

  For Address Family: IPv4 Unicast
   BGP neighbor version 2
   Update group: 0.2
   AF-dependant capabilities:
     Graceful Restart Capability advertised and received
       Neighbor preserved the forwarding state during latest restart
       Local restart time is 120, RIB purge time is 600 seconds
       Maximum stalepath time is 360 seconds
       Remote Restart time is 120 seconds
   Route refresh request: received 0, sent 0
   Policy for incoming advertisements is pass-all
   Policy for outgoing advertisements is pass-all
   1 accepted prefixes, 1 are bestpaths
   Prefix advertised 1, suppressed 0, withdrawn 0, maximum limit 524288
   Threshold for warning message 75%
   An EoR was not received during read-only mode

   Connections established 1; dropped 0
  Last reset 00:06:58, due to User clear requested (CEASE notification sent - administrative
reset)
   Time since last notification sent to neighbor: 00:06:58
   Error Code: administrative reset
   Notification data sent:
     None

 RP/0/RP0/CPU0:router# show bgp nei 5.5.5.2

 BGP neighbor is 5.5.5.2
  Remote AS 500, local AS 65000, external link
  Remote router ID 16.0.0.1
   BGP state = Established, up for 00:05:04
   BFD enabled (session up): mininterval: 318 multiplier: 4
   Last read 00:00:58, hold time is 180, keepalive interval is 60 seconds
   Precedence: internet
   Neighbor capabilities:
     Route refresh: advertised and received
     4-byte AS: advertised and received
     Address family IPv4 Unicast: advertised and received
   Received 8 messages, 0 notifications, 0 in queue
   Sent 9 messages, 1 notifications, 0 in queue
   Minimum time between advertisement runs is 30 seconds

  For Address Family: IPv4 Unicast
   BGP neighbor version 2
   Update group: 0.2
   AF-dependant capabilities:
     Graceful Restart Capability advertised and received
       Neighbor preserved the forwarding state during latest restart
```

```
      Local restart time is 120, RIB purge time is 600 seconds
      Maximum stalepath time is 360 seconds
      Remote Restart time is 120 seconds
   Route refresh request: received 0, sent 0
   Policy for incoming advertisements is pass-all
   Policy for outgoing advertisements is pass-all
   1 accepted prefixes, 0 are bestpaths
   Prefix advertised 1, suppressed 0, withdrawn 0, maximum limit 524288
   Threshold for warning message 75%
   An EoR was not received during read-only mode

   Connections established 1; dropped 0
  Last reset 00:07:01, due to User clear requested (CEASE notification sent - administrative
reset)
   Time since last notification sent to neighbor: 00:07:01
   Error Code: administrative reset
   Notification data sent:
     None

RP/0/RP0/CPU0:router# show bgp vrf one nei 3.12.1.2

BGP neighbor is 3.12.1.2, vrf one
 Remote AS 500, local AS 65000, external link
 Remote router ID 16.0.0.1
  BGP state = Established, up for 00:05:06
  BFD enabled (session up): mininterval: 119 multiplier: 10
  Last read 00:00:01, hold time is 180, keepalive interval is 60 seconds
  Precedence: internet
  Neighbor capabilities:
    Route refresh: advertised and received
    4-byte AS: advertised and received
    Address family IPv4 Unicast: advertised and received
  Received 9 messages, 0 notifications, 0 in queue
  Sent 9 messages, 1 notifications, 0 in queue
  Minimum time between advertisement runs is 0 seconds

 For Address Family: IPv4 Unicast
  BGP neighbor version 2
  Update group: 0.2
  AF-dependant capabilities:
    Graceful Restart Capability advertised and received
      Neighbor preserved the forwarding state during latest restart
      Local restart time is 120, RIB purge time is 600 seconds
      Maximum stalepath time is 360 seconds
      Remote Restart time is 120 seconds
  Route refresh request: received 0, sent 0
  Policy for incoming advertisements is pass-all
  Policy for outgoing advertisements is pass-all
  1 accepted prefixes, 1 are bestpaths
  Prefix advertised 0, suppressed 0, withdrawn 0, maximum limit 524288
  Threshold for warning message 75%
  An EoR was not received during read-only mode

  Connections established 1; dropped 0
 Last reset 00:07:04, due to User clear requested (CEASE notification sent - administrative
reset)
  Time since last notification sent to neighbor: 00:07:04
  Error Code: administrative reset
  Notification data sent:
    None
```

bestpath as-path multipath-relax (BGP IPv4 RT-filter address family)

To enable relaxed AS-path comparison for RT-filter multipath, use the **bestpath as-path multipath-relax** command in the IPv4 RT-constraint address family configuration mode. To disable this configuration, use the **no** form of this command.

bestpath as-path multipath-relax

This command has no keywords or arguments.

Command Default

None

Command History

Release	Modification
Release 26.3.1	This command was introduced.

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Task ID

Task ID	Operations
bgp	read, write

Examples

This example enables relaxed AS-path comparison for RT-filter multipath.

```
Router(config)# router bgp 140
Router(config-bgp)# address-family ipv4 rt-filter
Router(config-bgp-af)# bestpath as-path multipath-relax
```

To disable the configuration, use the **no bestpath as-path multipath-relax** command.

bgp as-path-loopcheck

To enable loop checking in the autonomous system path of the prefixes advertised by internal Border Gateway Protocol (iBGP) peers, use the **bgp as-path-loopcheck** command in an appropriate configuration mode. To restore the system to its default condition, use the **no** form of this command.

bgp as-path-loopcheck
no bgp as-path-loopcheck

Syntax Description

This command has no keywords or arguments.

Command Default

When you do not specify this command, loop checking is performed only for external peers.

Command Modes

Router configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

Task ID

Task ID	Operations
bgp	read, write

Examples

The following example shows how to configure an autonomous system path for loop checking iBGP peers:

```
RP/0/RP0/CPU0:router(config)# router bgp 6
RP/0/RP0/CPU0:router(config-bgp)# bgp as-path-loopcheck
```

bgp attribute-download

To enable Border Gateway Protocol (BGP) attribute download, use the **bgp attribute-download** command in an appropriate configuration mode. To disable BGP attribute download, use the **no** form of this command.

bgp attribute-download
no bgp attribute-download

Syntax Description

This command has no keywords or arguments.

Command Default

BGP attribute download is not enabled.

Command Modes

IPv4 unicast address family configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

When BGP attribute download is enabled using the **bgp attribute-download** command, BGP reinstalls all routes whose attributes are not currently in the RIB. Likewise, if the user disables BGP attribute download using the no form of the command, BGP reinstalls previously installed routes with a null key, and removes the attributes from the RIB.

Use the **bgp attribute-download** command to enable the Netflow BGP data export function. When attribute download is enabled, BGP downloads the attribute information for prefixes (community, extended community, and as-path) to the Routing Information Base (RIB) and Forwarding Information Base (FIB). This enables FIB to associate the prefixes with attributes and send the Netflow statistics along with the associated attributes.

Task ID

Task ID	Operations
bgp	read, write

Examples

The following example shows the BGP routes before and after BGP attribute download is enabled and shows how to enable BGP attribute download on BGP router 50:

```
RP/0/RP0/CPU0:router# show route bgp

B    100.0.1.0/24 [200/0] via 10.0.101.1, 00:00:37
B    100.0.2.0/24 [200/0] via 10.0.101.1, 00:00:37
B    100.0.3.0/24 [200/0] via 10.0.101.1, 00:00:37
B    100.0.4.0/24 [200/0] via 10.0.101.1, 00:00:37
B    100.0.5.0/24 [200/0] via 10.0.101.1, 00:00:37

RP/0/RP0/CPU0:router(config)# router bgp 50
RP/0/RP0/CPU0:router(config-bgp)# address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-bgp-af)# bgp attribute-download
!
!
!
```

```
RP/0/RP0/CPU0:router# show route bgp

B    100.0.1.0/24 [200/0] via 10.0.101.1, 00:00:01
      Attribute ID 0x2
B    100.0.2.0/24 [200/0] via 10.0.101.1, 00:00:01
      Attribute ID 0x2
B    100.0.3.0/24 [200/0] via 10.0.101.1, 00:00:01
      Attribute ID 0x2
B    100.0.4.0/24 [200/0] via 10.0.101.1, 00:00:01
      Attribute ID 0x2
B    100.0.5.0/24 [200/0] via 10.0.101.1, 00:00:01
      Attribute ID 0x2
```

bgp auto-policy-soft-reset disable

To disable an automatic soft reset of Border Gateway Protocol (BGP) peers when their configured route policy is modified, use the **bgp auto-policy-soft-reset disable** command in an appropriate configuration mode. To re-enable automatic soft reset of BGP peers, use the **no** form of this command.

bgp auto-policy-soft-reset disable
no bgp auto-policy-soft-reset disable

Syntax Description

This command has no keywords or arguments.

Command Default

Automatic soft reset of peers is enabled.

Command Modes

Router configuration

VRF configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

If the inbound policy changes, it is not always possible to perform a soft reset. This is the case if the neighbor does not support route refresh and soft-reconfiguration inbound is not configured for the neighbor. In such instances, a message is logged in the system log indicating that a manual hard reset is needed.

Task ID

Task ID	Operations
bgp	read, write

Examples

The following example shows how to disable an automatic soft reset of BGP peers when their configured route policy is modified:

```
RP/0/RP0/CPU0:router(config)# router bgp 6
RP/0/RP0/CPU0:router(config-bgp)# bgp auto-policy-soft-reset disable
```

bgp bestpath origin-as allow invalid

To permit all paths marked with an 'invalid' origin-as by RPKI to be considered for BGP best path computation, use the **bgp bestpath origin-as allow invalid** command in the router configuration mode. This configuration can also be made in the address family submode.

bgp bestpath origin-as allow invalid

Syntax Description

This command has no keywords or arguments.

Command Default

By default, prefixes marked with an 'invalid' origin-as are not considered for BGP best path computation when the router is performing origin-as validation.

Command Modes

Router configuration

Address family configuration

Command History

Release	Modification
Release 6.0	This command was introduced

Usage Guidelines

Configuring the **bgp bestpath origin-as allow invalid** command allows paths marked with an 'invalid' origin-as to be considered for best path computation. This can be limited to an address family by configuring it at the address-family submode.

This configuration takes effect only when the **bgp bestpath origin-as use validity** configuration is enabled.

Prior to Release 6.5.1, you can configure the **bgp bestpath origin-as use validity** command in default VRF only.

From Release 6.5.1 onwards, you can configure the **bgp bestpath origin-as use validity** command in any VRF.

Task ID

Task ID	Operation
bgp	read, write

Examples

The following example shows how to permit all invalid paths to be considered for BGP best-path selection:

```
Router#configure
Router(config)#router bgp 50000
Router(config-bgp)#bgp bestpath origin-as allow invalid
```

bgp bestpath origin-as use validity

To enable the BGP Origin AS Validation feature (RPKI) and allow the validity states of BGP paths to be taken into consideration in the bestpath process, use the **bgp bestpath origin-as use validity** command. This can be configured in router configuration mode and address family submode.

bgp bestpath origin-as use validity

Syntax Description

This command has no keywords or arguments.

Command Default

By default, the best path computation does not take RPKI states into account.

Command Modes

Router configuration

Address family configuration

Command History

Release	Modification
Release 6.0	This command was introduced

Usage Guidelines

There are three RPKI states - valid, invalid, and not found. When the **bgp bestpath origin-as use validity** command is configured, only paths marked with 'valid' or 'not found' are considered as best path candidates. When the **bgp bestpath origin-as allow invalid** command is configured, paths marked as 'invalid' are also considered but preference is given to routes marked 'valid' over those marked 'invalid'.

Prior to Release 6.5.1, you can configure the **bgp bestpath origin-as use validity** command in default VRF only.

From Release 6.5.1 onwards, you can configure the **bgp bestpath origin-as use validity** command in any VRF.

Task ID

Task ID	Operation
bgp	read, write

Examples

The following example shows how to enable the validity states of BGP paths to affect the path's preference when performing best-path selection:

```
Router#configure
Router(config)#router bgp 50000
Router(config-bgp)#bgp bestpath origin-as use validity
```

bgp bestpath as-path ignore

To ignore the autonomous system path length when calculating preferred paths, use the **bgp bestpath as-path ignore** command in an appropriate configuration mode. To return the software to the default state in which it considers the autonomous system path length when calculating preferred paths, use the **no** form of this command.

bgp bestpath as-path ignore
no bgp bestpath as-path ignore

Syntax Description

This command has no keywords or arguments.

Command Default

The autonomous system path length is used (not ignored) when a best path is selected.

Command Modes

Router configuration

VRF configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

Use the **bgp bestpath as-path ignore** command to ignore the length of autonomous system paths when the software selects a preferred path. When the best path is selected, if this command is specified, all steps are performed as usual except comparison of the autonomous path length between candidate paths.

Task ID

Task ID	Operations
bgp	read, write

Examples

The following example shows how to configure the software to ignore the autonomous system length when performing best-path selection:

```
RP/0/RP0/CPU0:router(config)# router bgp 65000
RP/0/RP0/CPU0:router(config-bgp)# bgp bestpath as-path ignore
```

bgp bestpath compare-routerid

To compare identical routes received from external BGP (eBGP) peers during the best-path selection process and select the route with the lowest router ID, use the **bgp bestpath compare-routerid** command in an appropriate configuration mode. To disable comparing identical routes received from eBGP peers during best-path selection, use the **no** form of this command.

bgp bestpath compare-routerid
no bgp bestpath compare-routerid

Syntax Description

This command has no keywords or arguments.

Command Default

The software does not select a new best path if it is the same as the current best path (according to the BGP selection algorithm) except for the router ID.

Command Modes

Router configuration

VRF configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

Use the **bgp bestpath compare-routerid** command to affect how the software selects the best path, in the case where there are two paths of equal cost according to the BGP selection algorithm. This command is used to force the software to select the path with the lower router ID as the best path. If this command is not used, the software continues to use whichever path is currently the best path, regardless of which has the lower router ID.

Task ID

Task ID	Operations
bgp	read, write

Examples

The following example shows how to configure the BGP speaker in autonomous system 500 to compare the router IDs of similar paths:

```
RP/0/RP0/CPU0:router(config)# router bgp 500
RP/0/RP0/CPU0:router(config-bgp)# bgp bestpath compare-routerid
```

bgp bestpath cost-community ignore

To configure a router that is running the Border Gateway Protocol (BGP) to not evaluate the cost community attribute during the best-path selection process, use the **bgp bestpath cost-community ignore** command in an appropriate configuration mode. To restore the system to its default condition, use the **no** form of this command.

bgp bestpath cost-community ignore
no bgp bestpath cost-community ignore

Syntax Description

This command has no keywords or arguments.

Command Default

The behavior of this command is enabled by default until the cost community attribute is manually configured.

Command Modes

Router configuration

VRF configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

Use the **bgp bestpath cost-community ignore** command to disable the evaluation of the cost community attribute to help isolate problems and troubleshoot issues that relate to BGP path selection. This command can also be used to delay the activation of cost community attribute evaluation so that cost community filtering can be deployed in a large network at the same time.

Task ID

Task ID	Operations
bgp	read, write

Examples

The following example shows how to configure a router to not evaluate the cost community attribute during the best-path selection process:

```
RP/0/RP0/CPU0:router(config)# router bgp 500
RP/0/RP0/CPU0:router(config-bgp)# bgp bestpath cost-community ignore
```

bgp bestpath med always

To allow the comparison of the Multi Exit Discriminator (MED) for paths from neighbors in different autonomous systems, use the **bgp bestpath med always** command in an appropriate configuration mode. To disable considering the MED attribute in comparing paths, use the **no** form of this command.

bgp bestpath med always
no bgp bestpath med always

Syntax Description

This command has no keywords or arguments.

Command Default

The software does not compare MEDs for paths from neighbors in different autonomous systems.

Command Modes

Router configuration

VRF configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

The MED is one of the parameters that is considered by the software when selecting the best path among many alternative paths. The software chooses the path with the lowest MED.

By default, during the best-path selection process, the software makes a MED comparison only among paths from the same autonomous system. This command changes the default behavior of the software by allowing comparison of MEDs among paths regardless of the autonomous system from which the paths are received.

When the **bgp bestpath med always** command is not enabled and distributed BGP is configured, speakers calculate partial best paths only (executes the best-path steps up to the MED comparison) and send them to BGP Routing Information Base (bRIB). bRIB calculates the final best path (executes all the steps in the best-path calculation). When the **bgp bestpath med always** command is enabled and distributed BGP is configured, speakers can compare the MED across all ASs, allowing the speaker to calculate a single best path to send it to bRIB. bRIB is the ultimate process that calculates the final best path, but when the **bgp bestpath med always** command is enabled, the speakers send a single best path instead of potentially sending multiple, partial best paths

Task ID

Task ID	Operations
bgp	read, write

Examples

The following example shows how to configure the Border Gateway Protocol (BGP) speaker in autonomous system 100 to compare MEDs among alternative paths, regardless of the autonomous system from which the paths are received:

```
RP/0/RP0/CPU0:router(config)# router bgp 100
```

```
RP/0/RP0/CPU0:router(config-bgp)# bgp bestpath med always
```

bgp bestpath med confed

To enable Multi Exit Discriminator (MED) comparison among paths learned from confederation peers, use the **bgp bestpath med confed** command in an appropriate configuration mode. To disable the software from considering the MED attribute in comparing paths, use the **no** form of this command.

bgp bestpath med confed
no bgp bestpath med confed

Syntax Description

This command has no keywords or arguments.

Command Default

The software does not compare the MED of paths containing only confederation segments, or paths containing confederation segments followed by an AS_SET, with the MED of any other paths.

Command Modes

Router configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

By default, the MED of the following paths is not compared with the MED of any other path:

- Paths with an empty autonomous system path
- Paths beginning with an AS_SET
- Paths containing only confederation segments
- Paths containing confederation segments followed by an AS_SET

Use the **bgp bestpath med confed** command to affect how the following types of paths are treated in the BGP best-path algorithm:

- Paths containing only confederation segments
- Paths containing confederation segments followed by an AS_SET

The MED for paths that start with an AS_SEQUENCE or that start with confederation segments followed by an AS_SEQUENCE only is compared with the MED of other paths that share the same first autonomous system number in the autonomous system sequence (the neighbor autonomous system number). This behavior is not affected by the **bgp bestpath med confed** command.

As an example, suppose that autonomous systems 65000, 65001, 65002, and 65004 are part of a confederation, but autonomous system 1 is not. Suppose that for a particular route, the following paths exist:

- Path 1: 65000 65004, med = 2, IGP metric = 20
- Path 2: 65001 65004, med = 3, IGP metric = 10
- Path 3: 65002 1, med = 1, IGP metric = 30

If the **bgp bestpath med confed** command is enabled, the software selects path 1 as the best path because it:

- Has a lower MED than path 2
- Has a lower IGP metric than path 3

The MED is not compared with path 3 because it has an external autonomous system number (that is, an AS_SEQUENCE) in the path. If the **bgp bestpath med confed** command is not enabled, then MED is not

compared between any of these paths. Consequently, the software selects path 2 as the best path because it has the lowest IGP metric.

Task ID

Task ID	Operations
bgp	read, write

Examples

The following command shows how to enable Border Gateway Protocol (BGP) software to compare MED values for paths learned from confederation peers:

```
RP/0/RP0/CPU0:router(config)# router bgp 210
RP/0/RP0/CPU0:router(config-bgp)# bgp bestpath med confed
```

bgp bestpath med missing-as-worst

To have the software consider a missing Multi Exit Discriminator (MED) attribute in a path as having a value of infinity, making the path without a MED value the least desirable path, use the **bgp bestpath med missing-as-worst** command in an appropriate configuration mode. To disable considering the MED attribute in comparing paths, use the **no** form of this command.

bgp bestpath med missing-as-worst
no bgp bestpath med missing-as-worst

Syntax Description

This command has no keywords or arguments.

Command Default

The software assigns a value of 0 to the missing MED, causing the path with the missing MED attribute to be considered as the best possible MED.

Command Modes

Router configuration

VRF configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

Task ID

Task ID	Operations
bgp	read, write

Examples

The following example shows how to direct the Border Gateway Protocol (BGP) software to consider a missing MED attribute in a path as having a value of infinity, making this path the least desirable path:

```
RP/0/RP0/CPU0:router(config)# router bgp 210
RP/0/RP0/CPU0:router(config-bgp)# bgp bestpath med missing-as-worst
```

bgp client-to-client reflection disable

To disable reflection of routes between route-reflection clients using a Border Gateway Protocol (BGP) route reflector, use the **bgp client-to-client reflection disable** command in address family configuration mode. To re-enable client-to-client reflection, use the **no** form of this command.

bgp client-to-client reflection [cluster-id *cluster-id]* **disable**
no bgp client-to-client reflection [cluster-id *cluster-id]* **disable**

Syntax Description

cluster-id *cluster-id*	(Optional) Cluster ID for which intra-cluster route reflection is to be disabled; maximum of 4 bytes. Cluster ID can be entered either as an IP address or value. Range is 1 to 4294967295.

Syntax Description

This command has no keywords or arguments.

Command Default

Client-to-client reflection is enabled.

Command Modes

Address family configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

By default, the clients of a route reflector are not required to be fully meshed and the routes from a client are reflected to other clients. However, if the clients are fully meshed, route reflection is not required.

By default, the clients of a route reflector that are part of the same cluster are not required to be fully meshed and the routes from a client are reflected to other clients. However, if the clients are fully meshed, route reflection is not required. If the cluster-id is not specified, then this command disables intra-cluster route reflection for all clusters.

Examples

In this example, the three neighbors are fully meshed, so client-to-client reflection is disabled:

```
RP/0/RP0/CPU0:router(config)# router bgp 65534
RP/0/RP0/CPU0:router(config-bgp)# bgp cluster-id 2
RP/0/RP0/CPU0:router(config-bgp)# address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-bgp-af)# bgp client-to-client reflection cluster-id 2 disable
RP/0/RP0/CPU0:router(config-bgp-af)# exit
RP/0/RP0/CPU0:router(config-bgp)# neighbor-group rrclients
RP/0/RP0/CPU0:router(config-bgp-nbrgrp)# remote-as 65534
RP/0/RP0/CPU0:router(config-bgp-nbrgrp)# bgp cluster-id 2
RP/0/RP0/CPU0:router(config-bgp-nbrgrp)# address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-bgp-nbrgrp-af)# route-reflector-client
RP/0/RP0/CPU0:router(config-bgp-nbrgrp-af)# exit
RP/0/RP0/CPU0:router(config-bgp-nbrgrp)# exit

RP/0/RP0/CPU0:router(config-bgp)# neighbor 192.168.253.21 use neighbor-group rrclients
RP/0/RP0/CPU0:router(config-bgp)# neighbor 192.168.253.22 use neighbor-group rrclients
RP/0/RP0/CPU0:router(config-bgp)# neighbor 192.168.253.23 use neighbor-group rrclients
```

Examples

In this example, the three neighbors are fully meshed, so client-to-client reflection is disabled:

```
RP/0/RP0/CPU0:router(config)# router bgp 65534
RP/0/RP0/CPU0:router(config-bgp)# address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-bgp-af)# bgp client-to-client reflection disable
RP/0/RP0/CPU0:router(config-bgp-af)# exit
RP/0/RP0/CPU0:router(config-bgp)# neighbor-group rrclients
RP/0/RP0/CPU0:router(config-bgp-nbrgrp)# remote-as 65534
RP/0/RP0/CPU0:router(config-bgp-nbrgrp)# address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-bgp-nbrgrp-af)# route-reflector-client
RP/0/RP0/CPU0:router(config-bgp-nbrgrp-af)# exit
RP/0/RP0/CPU0:router(config-bgp-nbrgrp)# exit

RP/0/RP0/CPU0:router(config-bgp)# neighbor 192.168.253.21 use neighbor-group rrclients
RP/0/RP0/CPU0:router(config-bgp)# neighbor 192.168.253.22 use neighbor-group rrclients
```

bgp cluster-id

To configure the cluster ID if the Border Gateway Protocol (BGP) cluster has more than one route reflector, use the **bgp cluster-id** command in an appropriate configuration mode. To remove the cluster ID, use the **no** form of this command.

bgp cluster-id *cluster-id*
no bgp cluster-id [*cluster-id*]

Syntax Description

cluster-id	Cluster ID of this router acting as a route reflector; maximum of 4 bytes. Cluster ID can be entered either as an IP address or value. Range is 1 to 4294967295.

Command Default

A cluster ID is not configured.

Command Modes

Router configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

Together, a route reflector and its clients form a *cluster*. A cluster of clients usually has a single route reflector. In such instances, the cluster is identified by the software as the router ID of the route reflector. To increase redundancy and avoid a single point of failure in the network, a cluster might have more than one route reflector. If it does, all route reflectors in the cluster must be configured with the same 4-byte cluster ID so that a route reflector can recognize updates from route reflectors in the same cluster.

A single route reflector can also support multiple clusters. Each cluster is identified by a unique cluster-id. The cluster-id configured by the **bgp cluster-id** command is taken as the default. If bgp cluster-id is not configured, the router ID for the default VRF identifies the default cluster. A neighbor can be associated with one cluster only, and the corresponding cluster-id is configured in neighbor configuration mode. If the cluster-id is not configured for a neighbor and the neighbor is a route reflector client, then the neighbor is assigned to the default cluster.

Task ID

Task ID	Operations
bgp	read, write

Examples

The following example shows how to configure the local router as one of the route reflectors serving the cluster. Neighbor 192.168.70.24 is assigned to the default cluster with cluster-id 1.

```
RP/0/RP0/CPU0:router(config)# router bgp 65534
RP/0/RP0/CPU0:router(config-bgp)# bgp cluster-id 1
RP/0/RP0/CPU0:router(config-bgp)# neighbor 192.168.70.24
RP/0/RP0/CPU0:router(config-bgp-nbr)# remote-as 65534
RP/0/RP0/CPU0:router(config-bgp-nbr)# address-family ipv4 unicast
```

```
RP/0/RP0/CPU0:router(config-bgp-nbr-af)# route-reflector-client
```

bgp confederation identifier

To specify a Border Gateway Protocol (BGP) confederation identifier, use the **bgp confederation identifier** command in an appropriate configuration mode. To remove the confederation identifier, use the **no** form of this command.

bgp confederation identifier *as-number*
no bgp confederation identifier [*as-number*]

Syntax Description

as-number Autonomous system (AS) number that internally includes multiple autonomous systems.

- Range for 2-byte Autonomous system numbers (ASNs) is 1 to 65535.
- Range for 4-byte Autonomous system numbers (ASNs) in asplain format is 1 to 4294967295.
- Range for 4-byte Autonomous system numbers (ASNs) is asdot format is 1.0 to 65535.65535.

Command Default

No confederation identifier is configured.

Command Modes

Router configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

One way to reduce the internal BGP (iBGP) mesh is to divide an autonomous system into multiple autonomous systems and group them into a single confederation. Each autonomous system is fully meshed within itself, and has a few connections to another autonomous system in the same confederation. Although the peers in different autonomous systems have external BGP (eBGP) sessions, they exchange routing information as if they are iBGP peers. Specifically, the confederation maintains the next hop and local preference information, and that allows you to retain a single Interior Gateway Protocol (IGP) for all autonomous systems. To the outside world, the confederation looks like a single autonomous system.

Use the **bgp confederation identifier** command to specify the autonomous system number for the confederation. This autonomous system number is used when BGP sessions are established with external peers in autonomous systems that are not part of the confederation.

Task ID

Task ID	Operations
bgp	read, write

Examples

The following example shows how to divide the autonomous system into autonomous systems 4001, 4002, 4003, 4004, 4005, 4006, and 4007 with the confederation identifier 5. Neighbor 10.2.3.4 is a router inside the confederation. Neighbor 172.20.16.6 is outside the routing domain confederation. To the outside world, there appears to be a single autonomous system with the number 5.

```
RP/0/RP0/CPU0:router(config)# router bgp 4001
RP/0/RP0/CPU0:router(config-bgp)# bgp confederation identifier 5
RP/0/RP0/CPU0:router(config-bgp)# bgp confederation peers 4002
RP/0/RP0/CPU0:router(config-bgp)# bgp confederation peers 4003
RP/0/RP0/CPU0:router(config-bgp)# bgp confederation peers 4004
RP/0/RP0/CPU0:router(config-bgp)# bgp confederation peers 4005
RP/0/RP0/CPU0:router(config-bgp)# bgp confederation peers 4006
RP/0/RP0/CPU0:router(config-bgp)# bgp confederation peers 4007
RP/0/RP0/CPU0:router(config-bgp)# neighbor 10.2.3.4
RP/0/RP0/CPU0:router(config-bgp-nbr)# remote-as 4002
RP/0/RP0/CPU0:router(config-bgp-nbr)# exit
RP/0/RP0/CPU0:router(config-bgp)# exit
RP/0/RP0/CPU0:router(config-bgp-nbr)# neighbor 172.20.16.6
RP/0/RP0/CPU0:router(config-bgp-nbr)# remote-as 4009
```

bgp confederation peers

To configure the autonomous systems that belong to the confederation, use the **bgp confederation peers** command in an appropriate configuration mode. To remove the autonomous system from the confederation, use the **no** form of this command.

bgp confederation peers [*as-number*]
no bgp confederation peers [*as-number*]

Syntax Description

as-number	Autonomous system (AS) numbers for Border Gateway Protocol (BGP) peers that belong to the confederation. • Range for 2-byte Autonomous system numbers (ASNs) is 1 to 65535. • Range for 4-byte Autonomous system numbers (ASNs) in asplain format is 1 to 4294967295. • Range for 4-byte Autonomous system numbers (ASNs) is asdot format is 1.0 to 65535.65535.

Command Default

No BGP peers are identified as belonging to the confederation.

Command Modes

Router configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

The autonomous systems specified in this command are visible internally to a confederation. Each autonomous system is fully meshed within itself. The bgp confederation identifier, on page 66 command specifies the confederation to which the autonomous systems belong.

To specify multiple autonomous systems, enter BGP confederation peer configuration mode then enter one *autonomous-system-number* for each command line.

Task ID

Task ID	Operations
bgp	read, write

Examples

The following example shows that autonomous systems 1090 and 1093 belong to a single confederation:

```
RP/0/RP0/CPU0:router(config)# router bgp 1090
RP/0/RP0/CPU0:router(config-bgp)# bgp confederation peers 1093
```

The following example shows that autonomous systems 1095, 1096, 1097, and 1098 belong to a single confederation:

```
RP/0/RP0/CPU0:router(config)# router bgp 1095
RP/0/RP0/CPU0:router(config-bgp)# bgp confederation peers
RP/0/RP0/CPU0:router(config-bgp-confed-peers)# 1096
RP/0/RP0/CPU0:router(config-bgp-confed-peers)# 1097
RP/0/RP0/CPU0:router(config-bgp-confed-peers)# 1098
```

bgp dampening

To enable Border Gateway Protocol (BGP) route dampening or change various BGP route dampening factors, use the **bgp dampening** command in an appropriate configuration mode. To disable route dampening and reset default values, use the **no** form of this command.

bgp dampening [*half-life* [*reuse suppress max-suppress-time*] | **route-policy** *route-policy-name*]
no bgp dampening [*half-life* [*reuse suppress max-suppress-time*] | **route-policy** *route-policy-name*]

Syntax Description

half-life	(Optional) Time (in minutes) after which a penalty is decreased. Once the route has been assigned a penalty, the penalty is decreased by half after the half-life period (which is 15 minutes by default). Penalty reduction happens every 5 seconds. Range of the half-life period is from 1 to 45 minutes.
reuse	(Optional) Value for route reuse if the flapping route penalty decreases and falls below the reuse value. When this happens, the route is unsuppressed. The process of unsuppressing routes occurs at 10-second increments. Range is 1 to 20000.
suppress	(Optional) Maximum penalty value. Suppress a route when its penalty exceeds the value specified. When this happens, the route is suppressed. Range is 1 to 20000.
max-suppress-time	(Optional) Maximum time (in minutes) a route can be suppressed. Range is 1 to 255. If the *half-life* value is allowed to default, the maximum suppress time defaults to 60 minutes.
route-policy *route-policy-name*	(Optional) Specifies the route policy to use to set dampening parameters.

Command Default

Route dampening is disabled.

half-life : 15 minutes

reuse : 750

suppress : 2000

max-suppress-time : four times *half-life* value

Command Modes

IPv4 address family configuration

IPv6 address family configuration

VPNv4 address family configuration

VRF IPv4 address family configuration

VPNv6 address family configuration

VRF IPv6 address family configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

Use the **bgp dampening** command without arguments to enable BGP route dampening with the default parameters. The parameters can be changed by setting them on the command line or specifying them with a routing policy.

Task ID

Task ID	Operations
bgp	read, write

Examples

The following example shows how to set the *half-life* value to 30 minutes, the *reuse* value to 1500, the *suppress* value to 10000, and the *max-suppress-time* to 120 minutes:

```
RP/0/RP0/CPU0:router(config)# router bgp 50
RP/0/RP0/CPU0:router(config-bgp)# address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-bgp-af)# bgp dampening 30 1500 10000 120
```

bgp default local-preference

To change the default local preference value, use the **bgp default local-preference** command in an appropriate configuration mode. To reset the local preference value to the default of 100, use the **no** form of this command.

bgp default local-preference *value*
no bgp default local-preference [*value*]

Syntax Description

value	Local preference value. Range is 0 to 4294967295. Higher values are preferable.

Command Default

Enabled with a value of 100.

Command Modes

Router configuration

VRF configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

Generally, the default value of 100 allows you to easily define a particular path as less preferable than paths with no local preference attribute. The preference is sent to all networking devices in the local autonomous system.

Task ID

Task ID	Operations
bgp	read, write

Examples

The following example shows how to raise the default local preference value from the default of 100 to 200:

```
RP/0/RP0/CPU0:router(config)# router bgp 200
RP/0/RP0/CPU0:router(config-bgp)# bgp default local-preference 200
```

bgp enforce-first-as disable

To disable the software from enforcing the first autonomous system path (known as the AS path) of a route received from an external Border Gateway Protocol (eBGP) peer to be the same as the configured remote autonomous system, use the **bgp enforce-first-as disable** command in an appropriate configuration mode. To re-enable enforcing the first AS path of a received route from an eBGP peer to be the same as the remote autonomous system, use the **no** form of this command.

bgp enforce-first-as disable
no bgp enforce-first-as disable

Syntax Description

This command has no keywords or arguments.

Command Default

By default, the software requires the first autonomous system (in the AS path) of a route received from an eBGP peer to be the same as the remote autonomous system configured.

Command Modes

Router configuration

VRF configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

By default, the software ignores any update received from an eBGP neighbor that does not have the autonomous system configured for that neighbor at the beginning of the AS path. When configured, the command applies to all eBGP peers of the router.

Task ID

Task ID	Operations
bgp	read, write

Examples

The following example shows a configuration in which incoming updates from eBGP neighbors are not checked to ensure the first AS number in the AS path is the same as the configured AS number for the neighbor:

```
RP/0/RP0/CPU0:router(config)# router bgp 100
RP/0/RP0/CPU0:router(config-bgp)# bgp enforce-first-as disable
```

bgp fast-external-fallover disable

To disable immediately resetting the Border Gateway Protocol (BGP) sessions of any directly adjacent external peers if the link used to reach them goes down, use the **bgp fast-external-fallover disable** command in an appropriate configuration mode. To disable this function and perform an immediate reset of BGP sessions when a link between peers is lost, use the **no** form of this command.

bgp fast-external-fallover disable
no bgp fast-external-fallover disable

Syntax Description

disable	Disables BGP fast external failover.

Command Default

BGP sessions of any directly adjacent external peers are immediately reset if the link used to reach them goes down.

Command Modes

Router configuration

VRF configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

By default, BGP sessions of any directly adjacent external peers are immediately reset, which allows the network to recover faster when links go down between BGP peers.

Task ID

Task ID	Operations
bgp	read, write

Examples

The following example shows how to disable the automatic resetting of BGP sessions:

```
RP/0/RP0/CPU0:router(config)# router bgp 109
RP/0/RP0/CPU0:router(config-bgp)# bgp fast-external-fallover disable
```

bgp graceful-restart

To enable graceful restart support, use the **bgp graceful-restart** command in an appropriate configuration mode. To disable this function, use the **no** form of this command.

bgp graceful-restart
no bgp graceful-restart

Syntax Description

This command has no keywords or arguments.

Command Default

Graceful restart support is not enabled.

Command Modes

Router configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

Use the **bgp graceful-restart** command to enable graceful restart functionality on the router, and also to advertise graceful restart to neighboring routers.

Note The **bgp graceful-restart** command with no options must be used to enable graceful restart before using the **bgp graceful-restart purge-time** , **bgp graceful-restart restart-time** , **bgp graceful-restart stalepath-time** , or **bgp graceful-restart graceful-reset** commands.

When graceful restart is enabled, the BGP graceful restart capability is negotiated with neighbors in the BGP OPEN message when the session is established. If the neighbor also advertises support for graceful restart, then graceful restart is activated for that neighbor session. If the neighbor does not advertise support for graceful restart, then graceful restart is not activated for that neighbor session even though it is enabled locally.

If you enter the **bgp graceful-restart** command after some BGP sessions are established, you must restart those sessions before graceful restart takes effect. Use the **clear bgp** command to restart sessions.

Unconfiguring a neighbor (**no neighbor** command) or BGP (**no router bgp** command) may trigger a graceful restart (GR) on the neighboring router. If you do not want the BGP to perform a graceful restart, disable the BGP neighbor by configuring an administrative shutdown under the neighbor settings before unconfiguring the neighbor or BGP.

Task ID

Task ID	Operations
bgp	read, write

Examples

The following example shows how to enable graceful restart:

```
RP/0/RP0/CPU0:router(config)#router bgp 3
```

```
RP/0/RP0/CPU0:router(config-bgp)#bgp graceful-restart
```

bgp graceful-restart graceful-reset

To invoke a graceful restart when configuration changes force a peer reset, use the **bgp graceful-restart graceful-reset** command in an appropriate configuration mode. To disable this function, use the **no** form of this command.

bgp graceful-restart graceful-reset
no bgp graceful-restart graceful-reset

Syntax Description

This command has no keywords or arguments.

Command Default

Graceful restart is not invoked when a configuration change forces a peer reset.

Command Modes

Router configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

BGP graceful restart must be enabled using the **bgp graceful-restart** command before enabling graceful reset using the **bgp graceful-restart graceful-reset** command.

Task ID

Task ID	Operations
bgp	read, write

Examples

The following example shows how to enable graceful reset:

```
RP/0/RP0/CPU0:router(config)#router bgp 3
RP/0/RP0/CPU0:router(config-bgp)# bgp graceful-restart graceful-reset
```

bgp graceful-restart purge-time

To specify the maximum time before stale routes are purged from the routing information base (RIB) when the local BGP process restarts, use the **bgp graceful-restart purge-time** command in an appropriate configuration mode. To set the purge timer time to its default value, use the **no** form of this command.

bgp graceful-restart purge-time *seconds*
no bgp graceful-restart purge-time *seconds*

Syntax Description

seconds	Maximum time before stale routes are purged. Time in seconds. Range is 0 to 6000.

Command Default

seconds : 600

Command Modes

Router configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

BGP graceful restart must be enabled using the **bgp graceful-restart** command before setting the purge time using the **bgp graceful-restart purge-time** command.

Task ID

Task ID	Operations
bgp	read, write

Examples

The following example shows how to change the BGP purge time to 800 seconds:

```
RP/0/RP0/CPU0:router(config)# router bgp 3
RP/0/RP0/CPU0:router(config-bgp)# bgp graceful-restart purge-time 800
```

bgp graceful-restart restart-time

To specify a user-predicted local BGP process maximum restart time, which is advertised to neighbors during session establishment, use the **bgp graceful-restart restart-time** command in an appropriate configuration mode. To set this restart time to its default value, use the **no** form of this command.

bgp graceful-restart restart-time *seconds*
no bgp graceful-restart restart-time *seconds*

Syntax Description

seconds	Maximum time advertised to neighbors. Time in seconds. Range is 1 to 4095.

Command Default

seconds : 120

Command Modes

Router configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

BGP graceful restart must be enabled using the **bgp graceful-restart** command before setting the restart timer using the **bgp graceful-restart restart-time** command.

Task ID

Task ID	Operations
bgp	read, write

Examples

The following example shows how to change the BGP graceful restart time to 400 seconds:

```
RP/0/RP0/CPU0:router(config)#router bgp 3
RP/0/RP0/CPU0:router(config-bgp)# bgp graceful-restart restart-time 400
```

bgp graceful-restart stalepath-time

To specify the maximum time to wait for an End-of-RIB message after a neighbor restarts, use the **bgp graceful-restart stalepath-time** command in an appropriate configuration mode. To set the stalepath timer time to its default value, use the **no** form of this command.

bgp graceful-restart stalepath-time *seconds*
no bgp graceful-restart stalepath-time *seconds*

Syntax Description

seconds	Maximum wait time. Time in seconds. Range is 1 to 4095.

Command Default

seconds : 360

Command Modes

Router configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

BGP graceful restart must be enabled using the **bgp graceful-restart** command before setting the stalepath time using the **bgp graceful-restart stalepath-time** command.

If the stalepath time is exceeded before an End-of-RIB message is received from a neighbor, paths learned from the neighbor are purged from the BGP routing table.

Task ID

Task ID	Operations
bgp	read, write

Examples

The following example shows how to change the stalepath time to 750 seconds:

```
RP/0/RP0/CPU0:router(config)# router bgp 3
RP/0/RP0/CPU0:router(config-bgp)# bgp graceful-restart stalepath-time 750
```

bgp import-delay

To enable delay for Border Gateway Protocol (BGP) batch import processing, use the **bgp import-delay** command in an appropriate configuration mode. To disable delay in batch import processing, use the no form of this command.

bgp import-delay *seconds milliseconds*
no bgp import-delay

Syntax Description

seconds	Specifies batch import processing delay in seconds. Range is 0 to 10 seconds.
milliseconds	Specifies batch import processing delay in milliseconds. Range is 0 to 999 seconds.

Command Default

No delay is configured.

Command Modes

Address-family VPNv4 Unicast

Address-family VPNv6 Unicast

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

Task ID

Task ID	Operation
bgp	read, write

This example shows how to set delay in batch import processing as two seconds and zero milliseconds:

```
RP/0/RP0/CPU0:router#configure
RP/0/RP0/CPU0:router(config)#router bgp 100
RP/0/RP0/CPU0:router(config-bgp)#address-family vpnv4 unicast
RP/0/RP0/CPU0:router(config-bgp-af)#bgp import-delay 2 0
```

bgp label-delay

To enable delay for Border Gateway Protocol (BGP) batch label processing, use the **bgp label-delay** command in an appropriate configuration mode. To disable delay in batch import processing, use the no form of this command.

bgp label-delay *seconds milliseconds*
no bgp label-delay

Syntax Description

seconds	Specifies batch label processing delay in seconds. Range is 0 to 10 seconds.
milliseconds	Specifies batch label processing delay in milliseconds. Range is 0 to 999 seconds.

Command Default

No delay is configured.

Command Modes

Address-family IPv4 Unicast

Address-family IPv6 Unicast

Address-family IPv4 Multicast

Address-family IPv6 Multicast

Address-family VPNv4 Unicast

Address-family VPNv6 Unicast

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

Task ID

Task ID	Operation
bgp	read, write

This example shows how to set delay in batch import processing as two seconds and zero milliseconds:

```
RP/0/RP0/CPU0:router#configure
RP/0/RP0/CPU0:router(config)#router bgp 100
RP/0/RP0/CPU0:router(config-bgp)#address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-bgp-af)#bgp label-delay 2 0
```

bgp lpts-secure-binding

To enable Local Packet Transport Services (LPTS) secure binding, use the **bgp lpts-secure-binding** command in BGP configuration mode. To disable the LPTS secure binding, use the **no** form of this command.

bgp lpts-secure-binding

Syntax Description

This command has no arguments or keywords.

Command Default

LPTS secure binding is not enabled.

Command Modes

Router BGP Configuration

Command History

Release	Modification
Release 7.10.1	This command was introduced.

Usage Guidelines

None.

Example

This example shows how to configure LPTS secure binding:

```
Router# router bgp 100
 Router(config-bgp)# bgp lpts-secure-binding
```

bgp log neighbor changes disable

To disable logging of Border Gateway Protocol (BGP) neighbor resets, use the **bgp log neighbor changes disable** command in an appropriate configuration mode. To re-enable logging of BGP neighbor resets, use the **no** form of this command.

bgp log neighbor changes disable
no bgp log neighbor changes disable

Syntax Description

This command has no keywords or arguments.

Command Default

BGP neighbor changes are logged.

Command Modes

Router configuration

VRF configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

Logging of BGP neighbor status changes (up or down) and resets is used for troubleshooting network connectivity problems and measuring network stability. Unexpected neighbor resets might indicate high error rates or high packet loss in the network, and should be investigated.

Status change message logging does not substantially affect performance, unlike, for example, enabling per-BGP update debugging. If the UNIX syslog facility is enabled, messages are sent by the software to the UNIX host running the syslog daemon so that the messages can be stored and archived on disk. If the UNIX syslog facility is not enabled, the status change messages are kept in the internal buffer of the router, and are not stored to disk.

The neighbor status change messages are not tracked if the **bgp log neighbor changes disable** command is disabled, except for the last reset reason, which is always available as output of the **show bgp neighbors** command.

Up and down messages for BGP neighbors are logged by the software by default. Use the **bgp log neighbor changes disable** command to stop logging BGP neighbor changes.

Task ID

Task ID	Operations
bgp	read, write

Examples

The following example shows how to prevent the logging of neighbor changes for BGP:

```
RP/0/RP0/CPU0:router(config)# router bgp 65530
RP/0/RP0/CPU0:router(config-bgp)# bgp log neighbor changes disable
```

bgp origin-as validation enable

To enable origin-AS validation, use the **bgp origin-as validation enable** command, in address-family configuration mode.

bgp origin-as validation enable

This command has no arguments or keywords.

Command Default

The BGP origin-AS validation disabled.

Command Modes

Router configuration

Address family configuration

Command History

Release	Modification
Release 6.6.1	This command was introduced

Usage Guidelines

None

Task ID

Task ID	Operation
bgp	read, write

Examples

The following example shows how to enable origin-AS validation in address-family configuration mode:

```
Router#configure
Router(config)#router bgp 50000
Router(config-bgp)#address-family ipv4 unicast
Router(config-bgp-af)#bgp origin-as validation enable
```

bgp origin-as validation signal ibgp

To signal origin-AS validity towards iBGP peers, use the **bgp origin-as validation signal ibgp** command, in Router configuration mode.

bgp origin-as validation signal ibgp

Syntax Description

ibgp	Specifies that the signal origin-AS validity is directed towards iBGP peers.

Command Default

By default, origin-AS validation is not signaled towards iBGP peers.

Command Modes

Router configuration

Address family configuration

Command History

Release	Modification
Release 6.6.1	This command was introduced

Usage Guidelines

None

Task ID

Task ID	Operation
bgp	read, write

Examples

The following example shows how to signal the validation of origin-AS towards iBGP peers in router configuration mode:

```
Router#configure
Router(config)#router bgp 50000
Router(config-bgp)#bgp origin-as validation signal ibgp
```

The following example shows how to signal the validation of origin-AS towards iBGP peers in address-family configuration mode:

```
Router#configure
Router(config)#router bgp 50000
Router(config-bgp)#address-family ipv4 unicast
Router(config-bgp-af)#bgp origin-as validation signal ibgp
```

bgp origin-as validation time

To set the origin-AS prefix validation time, use the **bgp origin-as validation time** command, in Router configuration mode.

bgp origin-as validation time {*prefix-validation-time* | **off**}

Syntax Description

prefix-validation-time	Sets the prefix validation time (in seconds). Range for prefix validation time is 5 to 60 seconds.
off	Sets off the automatic prefix validation after an RPKI update.

Command Default

The origin-AS prefix validation time is not set.

Command Modes

Router configuration

Address family configuration

Command History

Release	Modification
Release 6.6.1	This command was introduced

Usage Guidelines

None

Task ID

Task ID	Operation
bgp	read, write

Examples

The following example shows how to set the prefix validation time in router configuration mode:

```
Router#configure
Router(config)#router bgp 50000
Router(config-bgp)#bgp origin-as validation time 50
```

The following example shows how to set off the automatic prefix validation after an RPKI update in router configuration mode:

```
Router#configure
Router(config)#router bgp 50000
Router(config-bgp)#bgp origin-as validation time off
```

The following example shows how to set the prefix validation time in address-family configuration mode:

```
Router#configure
Router(config)#router bgp 50000
Router(config-bgp)#address-family ipv4 unicast
Router(config-bgp-af)#bgp origin-as validation time 50
```

The following example shows how to set off set off the automatic prefix validation after an RPKI update in address-family configuration mode:

```
Router#configure
Router(config)#router bgp 50000
Router(config-bgp)#address-family ipv4 unicast
Router(config-bgp-af)#bgp origin-as validation time off
```

bgp maximum neighbor

To control the maximum number of neighbors that can be configured on the router, use the **bgp maximum neighbor** command in an appropriate configuration mode. To set the neighbor limit to the default value, use the **no** form of this command.

bgp maximum neighbor *limit*
no maximum neighbor [*limit*]

Syntax Description

limit	Maximum number of neighbors. Range is 1 to 15000.

Command Default

Default limit is 10000

Command Modes

Router configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

Any attempt to configure the neighbor limit below 1 or above 15000 fails. Similarly, attempting to configure the limit below the number of neighbors currently configured fails. For example, if there are 3250 neighbors configured, you cannot set the *limit* below 3250.

Task ID

Task ID	Operations
bgp	write

Examples

The following example shows how to change the default maximum neighbor limit and set it to 1200:

```
RP/0/RP0/CPU0:router(config)#router bgp 65530
RP/0/RP0/CPU0:router(config-bgp)# bgp maximum neighbor 1200
```

bgp multipath as-path

To ignore as-path onwards while computing multipath, use the **bgp multipath as-path** command in XR Config mode.

bgp multipath as-path ignore onwards

Syntax Description

ignore	Ignores as-path related check for multipath selection.
onwards	Ignores everything as-path onwards for multipath selection.

Command Default

No default behavior or values

Command Modes

XR Config mode

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

When multiple connected routers start ignoring as-path onwards while computing multipath, it causes routing loops. Therefore, you should not configure the **bgp multipath as-path ignore onwards** command on routers that can form a loop.

Task ID

Task ID	Operations
bgp	read, write

Examples

This example shows how to ignore as-path while computing multipath.

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router bgp 100
RP/0/RP0/CPU0:router(config-bgp)# bgp multipath as-path ignore onwards
```

bgp policy accounting

To enable Border Gateway Protocol (BGP) policy accounting, use the **bgp policy accounting** command in interface configuration mode. To disable BGP policy accounting, use the **no** form of this command.

ipv4 bgp policy accounting { **input** | **output** { **destination-accounting** [**source-accounting**] | **source-accounting** [**destination-accounting**] } }

ipv6 bgp policy accounting { **input** | **output** { **destination-accounting** [**source-accounting**] | **source-accounting** [**destination-accounting**] } }

Syntax Description

input	Enables BGP policy accounting policy on the ingress IPv4 or IPv6 unicast interface.
destination-accounting	Enables accounting policy on the basis of the destination address.

Command Default

There is no BGP policy accounting.

Command Modes

Interface configuration

Command History

Release	Modification
Release 3.7.2	This command was introduced.

Usage Guidelines

- To use this command, you must be in a user group associated with a task group that includes the proper task IDs. If you suspect user group assignment is preventing you from using a command, contact your AAA administrator for assistance.
- To specify the accounting policy, the proper route policy configuration must be in place, matching specific BGP attributes using the **set traffic-index** command. In BGP router configuration mode, use the **table-policy** command to modify the accounting buckets when the IP routing table is updated with routes learned from BGP. To display accounting policy information, use the **show cef interface bgp-policy-statistics**, **show bgp policy**, and **show route bgp** commands.
- The functionality of the following keywords are not supported:
  - **output**
  - **source-accounting**
  - **output destination-accounting**
- This feature is applicable for the following address families:
  - IPv4
  - IPv6

Task ID

Task ID	Operations
network	read, write

Examples

The following example shows how to configure the BGP policy accounting:

For IPv4:

```
Router(config)# interface HundredGigE 0/5/0/0
Router(config-if)# ipv4 bgp policy accounting output source-accounting
Router(config-if)# commit
```

For IPv6:

```
Router(config)# interface HundredGigE 0/5/0/0
Router(config-if)# ipv6 bgp policy accounting output source-accounting
Router(config-if)# commit
```

bgp policy propagation input flow-tag

To match packets based on an incoming source, destination IP address or action (such as redirect, drop, PBTS) and redirect it to a specific VRF, use the **bgp policy propagation input flow-tag** command in the interface configuration mode.

bgp policy propagation input flow-tag { **destination** | **source** }

Syntax Description

bgp policy propagation input flow-tag	Enables flow-tag policy propagation on the specified interfaces.
destination	The packets are matched based on an incoming destination IP address and redirected to a specific VRF.
source	The packets are matched based on an incoming source IP address and redirect it to a specific VRF.

Command Default

None

Command Modes

Router configuration

Interface configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

Use this command to apply the flow-tag to a specified interface. The packets are matched based on an incoming source, destination IP address or action (such as redirect, drop, PBTS) and redirected to a specific VRF.

Note You will not be able to enable both QPPB and flow tag feature simultaneously on an interface.

Task ID

Task ID	Operation
bgp	read, write

bgp redistribute-internal

To allow the redistribution of internal Border Gateway Protocol (iBGP) routes into an Interior Gateway Protocol (IGP), such as Intermediate System-to-Intermediate System (IS-IS) or Open Shortest Path First (OSPF), use the **bgp redistribute-internal** command in an appropriate configuration mode. To disable the redistribution of iBGP routes into IGPs, use the **no** form of this command.

bgp redistribute-internal
no bgp redistribute-internal

Syntax Description

This command has no keywords or arguments.

Command Default

By default, iBGP routes are not redistributed into IGPs.

Command Modes

Router configuration

VRF configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

Use of the **bgp redistribute-internal** command requires the **clear route *** command to be issued to reinstall all BGP routes into the IP routing table.

Note Redistributing iBGP routes into IGPs may cause routing loops to form within an autonomous system. Use this command with caution.

Task ID

Task ID	Operations
bgp	read, write

Examples

The following example shows how to redistribute iBGP routes into OSPF:

```
RP/0/RP0/CPU0:router(config)#router bgp 1
RP/0/RP0/CPU0:router(config-bgp)# bgp redistribute-internal
RP/0/RP0/CPU0:router(config-bgp)# exit
RP/0/RP0/CPU0:router(config)# router ospf area1
RP/0/RP0/CPU0:router(config-router)# redistribute bgp 1
RP/0/RP0/CPU0:router(config-router)# end
RP/0/RP0/CPU0:router# clear route *
```

bgp router-id

To configure a fixed router ID for a Border Gateway Protocol (BGP)-speaking router, use the **bgp router-id** command in an appropriate configuration mode. To disable a fixed router ID, use the **no** form of this command.

bgp router-id *ip-address*
no bgp router-id [*ip-address*]

Syntax Description

ip-address	IP Version 4 (IPv4) address to use as the router ID. Normally, this should be an IPv4 address assigned to the router.

Command Default

If no router ID is configured in BGP, BGP attempts to use the global router ID if one is configured and available. Otherwise, BGP uses the highest IP address configured on a loopback interface.

Command Modes

Router configuration

VRF configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

If you do not use the **bgp router-id** command to configure a router ID, an IP address is not configured on any loopback interface, and no global router ID is configured, BGP neighbors remain down.

For more details on router IDs, see the *BGP Configuration Guide for Cisco NCS 5500 Series Routers*

Task ID

Task ID	Operations
bgp	read, write

Examples

The following example shows how to configure the local router with the router ID of 192.168.70.24:

```
RP/0/RP0/CPU0:router(config)# router bgp 100
RP/0/RP0/CPU0:router(config-bgp)#bgp router-id 192.168.70.24
```

bgp scan-time

To configure scanning intervals of Border Gateway Protocol (BGP)-speaking networking devices, use the **bgp scan-time** command in an appropriate configuration mode. To restore the scanning interval to its default value, use the **no** form of this command.

bgp scan-time *seconds*
no bgp scan-time
seconds

Syntax Description

seconds	Scanning interval (in seconds) of BGP routing information. Range is 5 to 3600 seconds.

Command Default

The default scanning interval is 60 seconds.

Command Modes

Router configuration

IPv4 address family configuration

IPv6 address family configuration

VPNv4 address family configuration

VPNv6 address family configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

Use the **bgp scan-time** command to change how frequently the software processes scanner tasks, such as conditional advertisement, dynamic MED changes, and periodic maintenance tasks.

Task ID

Task ID	Operations
bgp	read, write

Examples

The following example shows how to set the scanning interval for IPv4 unicast to 20 seconds:

```
RP/0/RP0/CPU0:router(config)# router bgp 64500
RP/0/RP0/CPU0:router(config-bgp)# address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-bgp-af)# bgp scan-time 20
```

This example shows how to set the scanning interval to 20 seconds:

```
RP/0/RP0/CPU0:router(config)# router bgp 64500
```

```
RP/0/RP0/CPU0:router(config-bgp-af)# bgp scan-time 20
```

bgp unsafe-ebgp-policy

To pass all routes on the eBGP neighbors, use the **bgp unsafe-ebgp-policy** command in router configuration mode.

bgp unsafe-ebgp-policy

Syntax Description

This command has no keywords or arguments.

Command Default

None

Command Modes

Router configuration

Command History

Release	Modification
Release 7.5.1	This command was introduced.

Usage Guidelines

Configuring the **bgp unsafe-ebgp-policy** command enables passing all routes on the eBGP neighbors without the need to configure route policies.

Task ID

Task ID	Operations
bgp	read, write

Examples

The following example shows how to enable the unsafe ebgp policy:

```
Router# configure
Router(config)# router bgp 65000
Router(config-bgp)# bgp unsafe-ebgp-policy
```

bgp update-delay

To set the maximum initial delay for a Border Gateway Protocol (BGP)-speaking router to send the first updates, use the **bgp update-delay** command in an appropriate configuration mode. To restore the initial delay to its default value, use the **no** form of this command.

bgp update-delay *seconds* [**always**]
nobgp update-delay [*seconds*][**always**]

Syntax Description

seconds	Delay in seconds for the router to send the first updates. Range is 0 to 3600.
always	(Optional) Specifies that the router always wait for the update delay time, even if all neighbors have finished sending their initial updates sooner.

Command Default

120 seconds

Command Modes

Router configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

When BGP is started, it waits a specified period of time for its neighbors to establish peering sessions and to complete sending their initial updates. After all neighbors complete their initial updates, or after the update delay timer expires, the best path is calculated for each route, and the software starts sending advertisements out to its peers. This behavior improves convergence time. If the software were to advertise a route as soon as it learned it, it would have to readvertise the route each time it learned a new path that was preferred over all previously learned paths.

Use the **bgp update-delay** command to tune the maximum time the software waits after the first neighbor is established until it starts calculating best paths and sending out advertisements.

Starting with Release 26.2.1, when **bgp update-delay always** is configured with a value greater than 2 minutes, and at least one dynamic neighbor is configured for Route Target Constraints (RTC) Address Family Identifier (AFI), Cisco IOS XR forces the RTC AFI out of read-only mode after 2 minutes instead of waiting for the full configured update-delay interval. The VPN AFI continues to follow the configured update-delay value and still depends on the RTC AFI coming out of read-only mode. This behavior allows the End-of-RIB (EOR) for RTC to be sent earlier, enabling peers to send VPN prefixes and EOR before the VPN AFI exits read-only mode. This change does not apply when the configured update-delay is 2 minutes or less.

Task ID

Task ID	Operations
bgp	read, write

Examples

The following example shows how to set the maximum initial delay to 240 seconds:

```
RP/0/RP0/CPU0:router(config)#router bgp 64530
RP/0/RP0/CPU0:router(config-bgp)# bgp update-delay 240
```

bgp write-limit

To modify the upper bounds on update message queue lengths or to enable desynchronization, use the **bgp write-limit** command in an appropriate configuration mode. To return the bounds to their default values and to disable desynchronization, use the **no** form of this command.

bgp write-limit *group-limit global-limit* [**desynchronize**]
no bgp write-limit [*group-limit global-limit*] [**desynchronize**]

Syntax Description

group-limit	Per-update group limit on the number of update messages the software queues. Range is 500 to 100000000. Group limit cannot be greater than the global limit.
global-limit	Global limit on the number of update messages the software queues. Range is 500 to 100000000.
desynchronize	(Optional) Enables desynchronization.

Command Default

group-limit : 50,000

global-limit : 250,000

Desynchronizationis off.

Command Modes

Router configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

Use the **bgp write-limit** command to configure both a per-update group and a global limit on the number of messages the software queues when updating peers. Increasing these limits can result in faster Border Gateway Protocol (BGP) convergence, but also may result in higher memory use during convergence. In addition, this command can be used to enable desynchronization. Desynchronization can decrease memory use and speed up convergence for the fastest neighbors if one or more neighbors in an update group process updates significantly slower than other neighbors in the same group. However, enabling desynchronization can cause a significant degradation in overall convergence time, especially if the router is experiencing high CPU utilization. For this reason, enabling desynchronization is discouraged.

Task ID

Task ID	Operations
bgp	read, write

Examples

The following example shows how to configure BGP to operate with a per-update group limit of 9000 messages and a global limit of 27,000 messages:

```
RP/0/RP0/CPU0:router(config)# router bgp 65000
```

```
RP/0/RP0/CPU0:router(config-bgp)#bgp write-limit 9000 27000
```

bmp-activate

To enable Border Gateway Protocol (BGP) Monitoring Protocol (BMP) logging for a neighbor, use the **bmp-activate server** command in neighbor configuration mode. To disable BMP logging for a neighbor, use the **no** form of this command.

bmp-activate server *server-id*

Syntax Description

server *server-id*	Enables monitoring by the BMP server specified by the *server-id* variable. You can configure multiple **bmp-activate** commands under same neighbor with different server IDs to enable monitoring by multiple BMP servers.

Command Default

No default behavior or values

Command Modes

Neighbor configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

Task ID

Task ID	Operations
bgp	read

Examples

The following example shows how to activate BMP on a neighbor with IP address 1.1.1.1, which is monitored by BMP server with server ID as 4:

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router bgp 100
RP/0/RP0/CPU0:router(config-bgp)# neighbor 1.1.1.1
RP/0/RP0/CPU0:router(config-bgp-nbr)# bmp-activate server 4
```

bmp server

To configure Border Gateway Protocol (BGP) Monitoring Protocol (BMP) server and to enter BMP server configuration mode, use the **bmp server** command in XR Config mode. To remove a particular BMP server configuration, use **no** form of this command.

bmp server {*server-id* | **all**}

Syntax Description

server-id	Specifies BMP server ID. Server ID range is 1 to 8.
all	Specifies all BMP servers.
description *LINE*	Specifies BMP server description. Description can be up to 250 alphanumeric characters.
dscp	Sets IP DiffServ CodePoint (DSCP). The DSCP value can be a number from 0 to 63, or it can be one of the following keywords: **default** , **ef** , **af11** , **af12** , **af13** , **af21** , **af22** , **af23** , **af31** , **af32** , **af33** , **af41** , **af42** , **af43** , **cs1** , **cs2** , **cs3** , **cs4** , **cs5** , **cs6** , or **cs7** .
host *host-name*	Specifies the hostname of BMP server. The hostname of the BMP server can be specified in IP address format (standard dot-decimal notation for IPv4 or colon-hexadecimal notation for IPv6) format or the string name which can be resolved into an IP address by the router.
initial-delay *delay-time*	Sets the delay, in seconds, before initial connect request is sent to a BMP server. The delay that you can set ranges from 1 to 3600 seconds. The default is 7 seconds.
initial-refresh {**delay** | **skip**}	Delay to initiate route refresh requests to BMP enabled neighbors. Configures the initial refresh options to handle refresh requests sent by the router to its BMP-enabled neighbors. Sets the delay, in seconds, before an initial refresh request is sent by the router to its BMP-enabled neighbors. The delay range is 1 to 3600 seconds with a default of 1 second. The default is not to skip refresh requests. Configures the router to skip sending any refresh requests to its BMP-enabled neighbors.

precedence	Sets the precedence values in the IP header. The precedence value can be a number from 0 to 7, or it can be one of the following keywords: **critical** —Set packets with critical precedence (5) **flash** — Set packets with flash precedence (3) **flash-override** —Set packets with flash override precedence (4) **immediate** —Set packets with immediate precedence (2) **internet** —Set packets with internetwork control precedence (6) **network** —Set packets with network control precedence (7) **priority** —Set packets with priority precedence (1) **routine** —Set packets with routine precedence (0) The default is internet (6) .
shutdown	Shuts down the TCP connection to BMP server.
stats-reporting-period	Specifies statistics reporting period, in seconds, to BMP servers. The reporting period that you can set ranges from 1 to 3600 seconds. The default is 0.
update-source *type interface-path-id*	Specifies the source (physical or virtual interface) to reach the BMP server. **Note** Use the **show interfaces** command to see a list of all interfaces currently configured on the router. For more information about the syntax for the router, use the question mark (?) online help function.
vrf *vrf-name*	Specifies VPN routing and forwarding (VRF) instance.

Command Default

For default values refer Syntax Description table.

Command Modes

XR Config mode

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

Task ID

Task ID	Operations
bgp	read

Examples

This example shows how to configure initial refresh delay of 30 seconds for BGP neighbors on BMP server with server ID as 4:

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# bmp server 4 initial-refresh delay 30
```

This example shows how to configure hostname of BMP server as 192.168.10.1:

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# bmp server 8 host 192.168.10.1 port 56
```

This example shows how to configure HundredGigE at location 0/0/0/1 as source interface to reach BMP server:

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# bmp server 5 update-source HundredGigE 0/0/0/1bmp server 5
update-source tenGigE 0/0/0/1
```

bind-source (rpki-server)

To specify a source interface for an RPKI session,, use the **bind-source** command in BGP rpki-server configuration mode. To remove the source interface for an RPKI session, use the **no** form of this command.

bind-source *interface*
no bind-source *interface*

Syntax Description

type	Interface type. For more information, use the question mark (?) online help function.
interface-path-id	Physical interface or virtual interface. **Note** Use the show interfaces command to see a list of all interfaces currently configured on the router. For more information about the syntax for the router, use the question mark (?) online help function.

Command Default

Bind source is not configured.

Command Modes

RPKI server configuration

Command History

Usage Guidelines

Use this command to specify the source address of the rpki session.

Task ID

Task ID	Operation
bgp	read, write

This example shows how to specify a source interface for an RPKI session and remove a source interface for an RPKI session:

```
RP/0/RP0/CPU0:router#configure
RP/0/RP0/CPU0:router(config)#router bgp 65000
RP/0/RP0/CPU0:router(config-bgp)#rpki server 192.0.2.0
RP/0/RP0/CPU0:router(config-bgp-rpki-server)# bind-source interface loopback 0
RP/0/RP0/CPU0:router(config-bgp-rpki-server)#commit
RP/0/RP0/CPU0:router(config-bgp-rpki-server)#no bind-source interface loopback 0
RP/0/RP0/CPU0:router(config-bgp-rpki-server)#commit
```

capability additional-paths receive

To advertise capability of receiving additional paths to the peer, use the **capability additional-paths receive** command in neighbor or neighbor-group or session-group configuration mode. To disable the capability of receiving additional paths, use the **no** form of this command.

capability additional-paths receive [**disable**]
no capability additional-paths receive

Syntax Description

disable	Disables advertising capability of receiving additional paths.

Command Default

Capability is disabled.

Command Modes

Neighbor configuration

Neighbor group configuration

Session group configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

Use the **capability additional-paths receive** command to selectively enable or disable additional paths receive capability negotiation for a particular neighbor or neighbor-group or session-group. Configuring **additional-paths receive** command in global address-family mode is a pre-requisite for negotiating additional paths receive capability with the peer.

If you enter the **capability additional-paths receive** command after some BGP sessions are established, you must restart those sessions for the new configuration to take effect. Use the **clear bgp** command to restart sessions.

Task ID

Task ID	Operation
bgp	read, write

The following example shows how to advertise capability of receiving additional paths:

```
RP/0/RP0/CPU0:router(config)#router bgp 100
RP/0/RP0/CPU0:router(config-bgp)#neighbor 10.2.3.4
RP/0/RP0/CPU0:router(config-bgp-nbr)#capability additional-paths receive
```

capability additional-paths send

To advertise capability of sending additional paths to the peer, use the **capability additional-paths send** command in neighbor or neighbor-group or session-group configuration mode. To disable the capability of sending additional paths, use the **no** form of this command.

capability additional paths send [**disable**]
no capability additional paths send

Syntax Description

disable	Disables advertise additional paths send capability

Command Default

Capability is disabled.

Command Modes

Neighbor configuration

Neighbor group configuration

Session group configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

Use the **capability additional paths send** command to selectively enable or disable additional paths send capability negotiation for a particular neighbor or neighbor-group or session-group. Configuring the **additional-paths send** command in global address-family mode is a pre-requisite for negotiating additional paths send capability with the peer.

You must restart the BGP sessions for the new configuration to take effect. Use the **clear bgp** command to restart sessions.

Task ID

Task ID	Operation
bgp	read, write

The following example shows how to advertise capability of sending additional paths to the peer:

```
RP/0/RP0/CPU0:router(config)# router bgp 100
RP/0/RP0/CPU0:router(config-bgp)# neighbor 10.2.3.4
RP/0/RP0/CPU0:router(config-bgp-nbr)# capability additional-paths send
```

capability orf prefix

To advertise prefix list-based Outbound Route Filter (ORF) capability to the Border Gateway Protocol (BGP) peer, use the **capability orf prefix** command in an appropriate configuration mode. To remove the **capability orf prefix** command from the configuration file and restore the system to its default condition in which the software does not advertise the capability, use the **no** form of this command.

capability orf prefix {**receive** | **send** | **both** | **none**}
no capability orf prefix [**receive** | **send** | **both** | **none**]

Syntax Description

receive	Sets the capability to receive the ORF from a specified neighbor.
send	Sets the capability to send the ORF to a specified neighbor.
both	Sets the capability to receive and send the ORF from or to a specified neighbor.
none	Sets the capability to no for ORF receive or send from or to a specified neighbor.

Command Default

The routing device does not receive or send route prefix filter lists.

Command Modes

IPv4 address family group configuration

IPv6 address family group configuration

IPv4 neighbor address family configuration

VRF neighbor IPv4 address family configuration

IPv4 neighbor group address family configuration

IPv6 neighbor group address family configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

The advertisement of the prefix list ORF capability by a BGP speaker indicates whether the speaker can send prefix lists to the specified neighbor and whether it accepts prefix lists from the neighbor. The speaker sends a prefix list if it indicated the ability to send them, and if the neighbor indicated it was willing to accept them. Similarly, the neighbor sends a prefix list to the speaker if it indicated the ability to send them and the speaker indicated the willingness to accept them.

Note The capability orf and prefix list filter specified by orf route-policy must be explicitly configured.

If the neighbor sends a prefix list and the speaker accepts it, the speaker applies the received prefix list, plus any locally configured outbound filters, to limit its outbound routing updates to the neighbor. Increased filtering prevents unwanted routing updates between neighbors and reduces resource requirements for routing update generation and processing.

Use the **capability orf prefix** command to set whether to advertise send and receive capabilities to the specified neighbor.

Note Sending a receive capability can adversely affect performance, because updates sent to that neighbor cannot be replicated for any other neighbors.

If this command is configured for a neighbor group or neighbor address family group, all neighbors using the group inherit the configuration. Values of commands configured specifically for a neighbor override inherited values.

Task ID

Task ID	Operations
bgp	read, write

Examples

The following example shows how to configure the **capability orf prefix** command:

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# route-policy orfqq
RP/0/RP0/CPU0:router:(config-rpl)# if orf prefix in (10.0.0.0/8 ge 20) then
RP/0/RP0/CPU0:router(config-rpl)# pass
RP/0/RP0/CPU0:router(config-rpl)# endif
RP/0/RP0/CPU0:router(config-rpl)# if orf prefix in (1910::16 ge 120) then
RP/0/RP0/CPU0:router(config-rpl)# pass
RP/0/RP0/CPU0:router(config-rpl)# endif
RP/0/RP0/CPU0:router(config-rpl)# end-policy
RP/0/RP0/CPU0:router(config)# router bgp 65530
RP/0/RP0/CPU0:router(config-bgp)# neighbor 10.0.101.1
RP/0/RP0/CPU0:router(config-bgp-nbr)# remote-as 65534
RP/0/RP0/CPU0:router(config-bgp-nbr)# address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-bgp-nbr-af)# route-policy pass-all out
RP/0/RP0/CPU0:router(config-bgp-nbr-af)# capability orf prefix both
RP/0/RP0/CPU0:router(config-bgp-nbr-af)# orf route-policy orfqq
```

capability suppress 4-byte-as

To suppress 4-byte AS capability from being advertised to the BGP peer, use the **capability suppress 4-byte-as** command in the appropriate configuration mode. To remove the **capability suppress 4-byte-as** command from the configuration and restore the system to the default condition, in which the software advertises the capability, use the **no** form of this command.

capability suppress 4-byte-as [**inheritance-disable**]
no capability suppress 4-byte-as

Syntax Description

inheritance-disable	Prevents capability suppress 4-type-as being inherited from the parent.

Command Default

4-byte-as capability is advertised to the BGP peer.

Command Modes

Neighbor configuration

Neighbor group configuration

Session group configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

By default, the software advertises the 4-byte AS capability to BGP peers. To override this default behavior, use the **capability suppress 4-byte-as** command under the command modes listed in the "Command Modes" section. If configured under the neighbor group or session group, all neighbors using the group inherit the configuration. Use the **no** option to remove the command.

Caution The BGP session resets automatically, if the 4-byte AS capability of an existing BGP session is changed by configuring **capability suppress 4-byte-as** or **capability suppress 4-byte-as inheritance-disable** .

Task ID

Task ID	Operations
bgp	read, write

Examples

The following example shows how to configure the **capability suppress 4-byte-as** command:

```
RP/0/RP0/CPU0:router# show bgp nei 10.3.3.3 conf
neighbor 10.3.3.3
 remote-as 65000             [n:internal]
 description PE3             []
 update-source Loopback0     [n:internal]
 address-family ipv4 unicast [n:internal]
```

```
RP/0/RP0/CPU0:router#show bgp nei 10.3.3.3
BGP neighbor is 10.3.3.3
 Remote AS 65000, local AS 65000, internal link
 Description: PE3
 Remote router ID 10.3.3.3
  BGP state = Established, up for 1w0d
  Last read 00:00:17, hold time is 180, keepalive interval is 60 seconds
  Precedence: internet
  Neighbor capabilities:
    Route refresh: advertised and received
    4-byte AS: advertised and received
    Address family IPv4 Unicast: advertised and received
  Received 25962 messages, 0 notifications, 0 in queue
  Sent 25968 messages, 1 notifications, 0 in queue
  Minimum time between advertisement runs is 0 seconds

 For Address Family: IPv4 Unicast
  BGP neighbor version 1
  Update group: 0.3
  Route refresh request: received 0, sent 0
  0 accepted prefixes, 0 are bestpaths
  Prefix advertised 0, suppressed 0, withdrawn 0, maximum limit 524288
  Threshold for warning message 75%
  An EoR was received during read-only mode

  Connections established 2; dropped 1
  Last reset 1w0d, due to BGP Notification sent: hold time expired
  Time since last notification sent to neighbor: 1w0d
  Error Code: hold time expired
  Notification data sent: None

RP/0/RP0/CPU0:router(config)#router bgp 65000
RP/0/RP0/CPU0:router(config-bgp)#neighbor 10.3.3.3
RP/0/RP0/CPU0:router(config-bgp-nbr)#capability suppress 4-byte-as
RP/0/RP0/CPU0:router(config-bgp-nbr)#commit
RP/0/RP0/CPU0:router(config-bgp-nbr)#end

RP/0/RP0/CPU0:router# show bgp nei 10.3.3.3

BGP neighbor is 10.3.3.3
 Remote AS 65000, local AS 65000, internal link
 Description: PE3
 Remote router ID 10.3.3.3
  BGP state = Established, up for 00:00:16
  Last read 00:00:11, hold time is 180, keepalive interval is 60 seconds
  Precedence: internet
  Neighbor capabilities:
    Route refresh: advertised and received
    Address family IPv4 Unicast: advertised and received
  Capability 4-byte-as suppress is configured
  Received 25966 messages, 0 notifications, 0 in queue
  Sent 25972 messages, 1 notifications, 0 in queue
  Minimum time between advertisement runs is 0 seconds

 For Address Family: IPv4 Unicast
  BGP neighbor version 1
  Update group: 0.2
  Route refresh request: received 0, sent 0
  0 accepted prefixes, 0 are bestpaths
  Prefix advertised 0, suppressed 0, withdrawn 0, maximum limit 524288
  Threshold for warning message 75%
  An EoR was received during read-only mode
```

```
  Connections established 3; dropped 2
  Last reset 00:00:43, due to Capabilty 4-byte-as configuration changed
  Time since last notification sent to neighbor: 1w0d
  Error Code: hold time expired
  Notification data sent: None
```

With the **inheritance-disable** keyword:

```
RP/0/RP0/CPU0:router(config-bgp)# neighbor 10.0.101.1
RP/0/RP0/CPU0:router(config-bgp-nbr)# capability suppress 4-byte-as inheritance-disable

RP/0/RP0/CPU0:router# show bgp neighbor 10.0.101.1 config
neighbor 10.0.101.1
 remote-as 1                 []
 address-family ipv4 unicast []

RP/0/RP0/CPU0:router# show bgp neighbor 10.0.101.1
BGP neighbor is 10.0.101.1
 Remote AS 1, local AS 100, external link
 Remote router ID 0.0.0.0
  BGP state = Idle
  Last read 00:00:00, hold time is 180, keepalive interval is 60 seconds
  Precedence: internet
  Received 0 messages, 0 notifications, 0 in queue
  Sent 0 messages, 0 notifications, 0 in queue
  Minimum time between advertisement runs is 30 seconds
```

cef load-balancing recursive oor mode dampening-and-dlb

To enable FIB dampening and Destination-based Load Balancing (DLB) features for avoiding ECMP Out-of-Resource (OOR) conditions, use the **cef load-balancing recursive oor mode dampening-and-dlb** command in global configuration mode.

cef load-balancing recursive oor mode dampening-and-dlb [**max-duration** *value* | **dampening resource-threshold** *mbb_threshold* **max-duration** *value* **dlb resource-threshold** *dlb_threshold*]

Syntax Description

dampening resource-threshold *mbb_threshold*	Specifies the FIB dampening threshold value. The value ranges from 1 through 99.
max-duration *value*	Specifies the FIB dampening maximum switchover duration. The range is from 1 second to 600 seconds.
dlb resource-threshold *dlb_threshold*	Specifies the DLB threshold value. The value ranges from 1 through 99.

Command Default

FIB dampening and DLB features are disabled.

Command Modes

Global configuration

Command History

Release	Modification
Release 24.2.1	This command was introduced.

Usage Guidelines

To enable FIB dampening and DLB features with default values, use the **cef load-balancing recursive oor mode dampening-and-dlb** command. The default hardware usage values for FIB dampening and DLB are 70% and 90% respectively. The default FIB dampening switchover duration is 300 seconds.

To manually configure the FIB dampening switchover duration, use the **cef load-balancing recursive oor mode dampening-and-dlb max-duration** *value* command.

To manually configure the FIB dampening threshold value, FIB dampening maximum switchover duration, and DLB threshold value, use the **cef load-balancing recursive oor mode dampening-and-dlb dampening resource-threshold** *mbb_threshold* **max-duration** *value* **dlb resource-threshold** *dlb_threshold* command.

Task ID

Task ID	Operations
cef	read, write

The following example shows how to enable FIB dampening and DLB features for avoiding ECMP OOR conditions.

```
Router(config)# cef load-balancing recursive oor mode dampening-and-dlb dampening
resource-threshold 40 max-duration 300 dlb resource-threshold 85
```

clear bgp

To reset a group of Border Gateway Protocol (BGP) neighbors, use the **clear bgp** command in XR EXEC mode.

clear bgp [**ipv4** {**unicast** | **multicast** | **labeled-unicast** | **all** | **tunnel** | **mdt**} | **ipv6** {**unicast** | **multicast** | **all** | **labeled-unicast**} | **all** {**unicast** | **multicast** | **all** | **labeled-unicast** | **mdt** | **tunnel**} | **vpnv4 unicast** | **vrf** {**vrf-name** | **all**} {**ipv4** {**unicast** | **labeled-unicast**} | **ipv6 unicast**} | **vpnv6 unicast**]

Syntax Description

ipv4	(Optional) Specifies IP Version 4 address prefixes.
unicast	(Optional) Specifies unicast address prefixes.
multicast	(Optional) Specifies multicast address prefixes.
labeled-unicast	(Optional) Specifies labeled unicast address prefixes.
all	(Optional) For subaddress families, specifies prefixes for all subaddress families.
tunnel	(Optional) Specifies tunnel address prefixes.
ipv6	(Optional) Specifies IP Version 6 address prefixes.
all	(Optional) For address family, specifies prefixes for all address families.
vpnv4 unicast	(Optional) Specifies VPNv4 unicast address families.
vrf	(Optional) Specifies VPN routing and forwarding (VRF).
vrf-name	Name of a VRF.
all	(Optional) For VRF, specifies all VRFs.
ipv4 { **unicast** | **labeled-unicast** }	(Optional) For VRF, specifies IPv4 unicast and labeled-unicast address families.
ipv6 unicast	(Optional) For VRF, specifies IPv6 unicast address prefixes.

Command Default

No default behavior or values

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

Use the **clear bgp** command to reset the sessions of the specified group of neighbors (hard reset); it removes the TCP connection to the neighbor, removes all routes received from the neighbor from the BGP table, and then re-establishes the session with the neighbor.

If the **graceful** keyword is specified, the routes from the neighbor are not removed from the BGP table immediately, but are marked as stale. After the session is re-established, any stale route that has not been received again from the neighbor is removed.

Task ID

Task ID	Operations
bgp	execute

Examples

The following example shows how to hard reset neighbor 10.0.0.1:

```
RP/0/RP0/CPU0:router# clear bgp 10.0.0.1
```

clear bgp dampening

To clear Border Gateway Protocol (BGP) route dampening information and unsuppress the suppressed routes, use the **clear bgp dampening** command in XR EXEC mode.

clear bgp dampening

Syntax Description

ipv4	Specifies IP Version 4 address prefixes.
unicast	Specifies unicast address prefixes.
multicast	Specifies multicast address prefixes.
labeled-unicast	Specifies labeled unicast address prefixes.
all	For subaddress families, specifies prefixes for all subaddress families.
ipv6	Specifies IP Version 6 address prefixes.
all	For address family, specifies prefixes for all address families.
vpnv4 unicast	Specifies VPNv4 unicast address families.
vrf	Specifies VPN routing and forwarding (VRF).
vrf-name	Name of a VRF.
all	For VRF, specifies all VRFs.
ipv4 { **unicast** | **labeled-unicast** }	For VRF, specifies IPv4 unicast and labeled-unicast address families.
ipv6 unicast	For VRF, specifies IPv6 unicast address families.
vpnv6 unicast	Specifies VPNv6 unicast address families.
ip-address	(Optional) IP address of the network about which to clear dampening information.
/mask-length	(Optional) Network mask applied to the IP address.

Command Default

If no IP address is specified, dampening information for all routes is cleared.

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

Always use the **clear bgp dampening** command for an individual address-family. The **all** option for address-families with clear bgp dampening should never be used during normal functioning of the system. For example, use

```
clear bgp ipv4 unicast dampening prefix x.x.x./y
```

Task ID

Task ID	Operations
bgp	execute

Examples

The following example shows how to clear the route dampening information for all 172.20.0.0/16 IPv4 unicast paths:

```
RP/0/RP0/CPU0:router# clear bgp ipv4 unicast dampening 172.20.0.0/16
```

clear bgp external

To clear all Border Gateway Protocol (BGP) external peers, use the **clear bgp external** command in XR EXEC mode.

clear bgp external

Syntax Description

ipv4	(Optional) Specifies IP Version 4 address prefixes.
unicast	(Optional) Specifies unicast address prefixes.
multicast	(Optional) Specifies multicast address prefixes.
labeled-unicast	(Optional) Specifies labeled unicast address prefixes.
all	(Optional) For subaddress families, specifies prefixes for all subaddress families.
ipv6	(Optional) Specifies IP Version 6 address prefixes.
all	(Optional) For address family, specifies prefixes for all address families.
vpnv4 unicast	(Optional) Specifies VPNv4 unicast address families.
vrf	(Optional) Specifies VPN routing and forwarding (VRF).
vrf-name	(Optional) Name of a VRF.
all	(Optional) For VRF, specifies all VRFs.
ipv4 { **unicast** | **labeled-unicast** }	(Optional) For VRF, specifies IPv4 unicast or labeled-unicast address families.
ipv6 unicast	(Optional) For VRF, specifies IPv6 unicast address families.
vpnv6 unicast	(Optional) Specifies VPNv6 unicast address families.
graceful	(Optional) Clears all external peers with a hard reset and a graceful restart. This option is available when an address family is not specified.

Command Default

No default behavior or value

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

Task ID

Task ID	Operations
bgp	execute

Examples

The following example shows how to clear all BGP external peers:

```
RP/0/RP0/CPU0:router# clear bgp external
```

clear bgp flap-statistics

To clear Border Gateway Protocol (BGP) flap counts for a specified group of routes, use the **clear bgp flap-statistics** command in XR EXEC mode.

clear bgp flap-statistics

Syntax Description

ipv4	Specifies IP Version 4 address prefixes.
unicast	Specifies unicast address prefixes.
multicast	Specifies multicast address prefixes.
labeled-unicast	Specifies labeled unicast address prefixes.
all	For subaddress families, specifies prefixes for all subaddress families.
ipv6	Specifies IP Version 6 address prefixes.
all	For address family, specifies prefixes for all address families.
vpnv4 unicast	Specifies VPNv4 unicast address families.
vrf	Specifies VPN routing and forwarding (VRF).
vrf-name	Name of a VRF.
all	For VRF, specifies all VRFs.
ipv4 { **unicast** | **labeled-unicast** }	For VRF, specifies IPv4 unicast or labeled-unicast address families.
ipv6 unicast	For VRF, specifies IPv6 unicast address families.
vpnv6 unicast	Specifies VPNv6 unicast address families.
regexp *regexp*	(Optional) Clears flap statistics for routes whose AS paths match the regular expression.
route-policy *route-policy-name*	(Optional) Clears flap statistics for the specific route policy.
network	(Optional) Network for which flap counts are to be cleared.
/mask-length	(Optional) Network mask of the network for which flap counts are to be cleared.
ip-address	(Optional) Neighbor address. Clears only flap statistics for routes received from this neighbor.

Command Default

No default behavior or value

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

Task ID

Task ID	Operations
bgp	execute

Examples

The following example shows how to clear the flap count for all routes (in all address families) originating in autonomous system 1:

```
RP/0/RP0/CPU0:router#clear bgp all all flap-statistics regexp _1$
```

The following example shows how to clear the flap count for all IPv4 unicast routes received from neighbor 172.20.1.1:

```
RP/0/RP0/CPU0:router# clear bgp ipv4 unicast flap-statistics 172.20.1.1
```

clear bgp nexthop performance-statistics

To reset the number of received notifications and the cumulative processing time for the Border Gateway Protocol (BGP) next hop, use the **clear bgp nexthop performance-statistics** command in XR EXEC mode.

clear bgp nexthop performance-statistics

Syntax Description

ipv4	Specifies IP Version 4 address prefixes.
unicast	Specifies unicast address prefixes.
multicast	Specifies multicast address prefixes.
labeled-unicast	Specifies labeled unicast address prefixes.
all	For subaddress families, specifies prefixes for all subaddress families.
tunnel	Specifies tunnel address prefixes.
mdt	Specifies IPv4 multicast distribution tree (MDT) address prefixes.
ipv6	Specifies IP Version 6 address prefixes.
all	For address family, specifies prefixes for all address families.
vpnv4 unicast	Specifies VPNv4 unicast address families.
vrf	Specifies VPN routing and forwarding (VRF).
vrf-name	Name of a VRF.
all	For VRF, specifies all VRFs.
ipv4 { **unicast** | **labeled-unicast** }	For VRF, specifies IPv4 unicast or labeled-unicast address families.
ipv6 unicast	For VRF, specifies IPv6 unicast address families.
vpnv6 unicast	Specifies VPNv6 unicast address families.

Command Default

No default behavior or values

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

Use the **clear bgp nexthop performance-statistics** command to reset the total number of notifications received from the Routing Information Base (RIB) and the cumulative next-hop processing time. The following information is cleared from the **show bgp nexthops** command output:

- Total critical notifications received

- Total noncritical notifications received
- Best path deleted after last walk
- Best path changed after last walk
- Next-hop table total number of critical and noncritical notifications (Notf) and the time of the last notification received from the RIB (LastRIB) columns (only entries that have a status of unreachable [UR])

Task ID

Task ID	Operations
bgp	execute

Examples

The following example shows how to clear next-hop performance statistics:

```
RP/0/RP0/CPU0:router# clear bgp vrf vrf_A nexthop performance statistics
```

clear bgp nexthop registration

To reregister a specified next hop with the Routing Information Base (RIB), use the **clear bgp nexthop registration** command in XR EXEC mode.

clear bgp nexthop registration nexthop-address *nexthop-address*

Syntax Description

ipv4	Specifies IP Version 4 address prefixes.
unicast	Specifies unicast address prefixes.
multicast	Specifies multicast address prefixes.
labeled-unicast	Specifies labeled-unicast address prefixes.
all	For subaddress families, specifies prefixes for all subaddress families.
tunnel	Specifies tunnel address prefixes.
mdt	Specifies IPv4 multicast distribution tree (MDT) address prefixes.
ipv6	Specifies IP Version 6 address prefixes.
all	For address family, specifies prefixes for all address families.
vpnv4 unicast	Specifies VPNv4 unicast address families.
vrf	Specifies VPN routing and forwarding (VRF).
vrf-name	Name of a VRF.
all	For VRF, specifies all VRFs.
ipv4 { **unicast** | **labeled-unicast** }	For VRF, specifies IPv4 unicast or labeled-unicast address families.
ipv6 unicast	For VRF, specifies IPv6 unicast address families.
vpnv6 unicast	Specifies VPNv6 unicast address families.
nexthop-address	Address of the next hop.

Command Default

No default behavior or values

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

Use the **clear bgp nexthop registration** command to perform an asynchronous registration of the next hop with the RIB. The **show bgp nexthops** command output shows a critical notification as the LastRIBEvent for the next hop when the **clear bgp nexthop registration** command is used.

Task ID

Task ID	Operations
bgp	execute

Examples

The following example shows how to reregister the next hop with the RIB:

```
RP/0/RP0/CPU0:router# clear bgp nexthop registration 10.1.1.1
```

clear bgp peer-drops

To clear the connection-dropped counter, use the **clear bgp peer-drops** command in XR EXEC mode.

clear bgp peer-drops {*****|*ip-address*}

Syntax Description

*****	Specifies all BGP neighbors.
ip-address	IP address of a specific network neighbor.

Command Default

No default behavior or values

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

Task ID

Task ID	Operations
bgp	execute

Examples

The following example shows how to clear the connection-dropped counter for all BGP neighbors:

```
RP/0/RP0/CPU0:router# clear bgp peer-drops *
```

clear bgp performance-statistics

To clear the performance statistics for all address families, use the **clear bgp performance-statistics** command.

clear bgp [**vrf** {*vrf-name* | **all**}] **performance-statistics**

Syntax Description

vrf	Specifies VPN routing and forwarding (VRF).
vrf-name	Name of a VRF.
all	For VRF, specifies all VRFs.

Syntax Description

This command has no keywords or arguments.

Command Default

No default behavior or values

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

Task ID

Task ID	Operations
bgp	execute

Examples

The following example shows how to clear the performance statistics for all address families:

```
RP/0/RP0/CPU0:router# clear bgp performance-statistics
```

clear bgp self-originated

To clear Border Gateway Protocol (BGP) routes that are self-originated, use the **clear bgp self-originated** command in XR EXEC mode.

clear bgp {**ipv4** {**unicast** | **multicast** | **labeled-unicast** | **all**} | **ipv6** {**unicast** | **multicast** | **labeled-unicast** | **all** } | **all** {**unicast** | **multicast** | **labeled-unicast** | **all**} | **vpnv4 unicast** | **vrf** {*vrf-name* | **all**} | **vpnv6 unicast**} **self-originated**

Syntax Description

ipv4	Specifies IP Version 4 address prefixes.
unicast	Specifies unicast address prefixes.
multicast	Specifies multicast address prefixes.
labeled-unicast	Specifies labeled unicast address prefixes.
all	For subaddress families, specifies prefixes for all subaddress families.
ipv6	Specifies IP Version 6 address prefixes.
all	For address family, specifies prefixes for all address families.
vpnv4 unicast	Specifies VPNv4 unicast address families.
vrf	Specifies VPN routing and forwarding (VRF).
vrf-name	Name of a VRF.
all	For VRF, specifies all VRFs.
ipv4 { **unicast** | **labeled-unicast** }	For VRF, specifies IPv4 unicast or labeled-unicast address families.
ipv6 unicast	For VRF, specifies IPv6 unicast address families.
vpnv6 unicast	Specifies VPNv6 unicast address families.

Command Default

No default behavior or values

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

Self-originated routes are routes locally originated by the **network** command, **redistribute** command, or **aggregate-address** command.

Task ID

Task ID	Operations
bgp	execute

Examples

The following example shows how to clear self-originated IPv4 routes:

```
RP/0/RP0/CPU0:router# clear bgp ipv4 unicast self-originated
```

clear bgp shutdown

To clear all Border Gateway Protocol (BGP) neighbors that shut down due to low memory, use the **clear bgp shutdown** command in XR EXEC mode.

clear bgp {**ipv4** {**unicast** | **multicast** | **labeled-unicast** | **all**} | **ipv6** {**unicast** | **multicast** | **labeled-unicast** | **all** } | **all** {**unicast** | **multicast** | **labeled-unicast** | **all**} | **vpnv4 unicast** | **vrf** {*vrf-name* | **all**} | **vpnv6 unicast**} **shutdown**

Syntax Description

ipv4	Specifies IP Version 4 address prefixes.
unicast	Specifies unicast address prefixes.
multicast	Specifies multicast address prefixes.
labeled-unicast	Specifies labeled unicast address prefixes.
all	For subaddress families, specifies prefixes for all subaddress families.
ipv6	Specifies IP Version 6 address prefixes.
all	For address family, specifies prefixes for all address families.
vpnv4 unicast	Specifies VPNv4 unicast address families.
vrf	Specifies VPN routing and forwarding (VRF).
vrf-name	Name of a VRF.
all	For VRF, specifies all VRFs.
ipv4 { **unicast** | **labeled-unicast** }	For VRF, specifies IPv4 unicast or labeled-unicast address families.
ipv6 unicast	For VRF, specifies IPv6 unicast address families.
vpnv6 unicast	Specifies VPNv6 unicast address families.

Command Default

No default behavior or values

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

Task ID

Task ID	Operations
bgp	execute

Examples

The following example shows how to clear all shut-down BGP neighbors:

```
RP/0/RP0/CPU0:router# clear bgp shutdown
```

clear bgp soft

To soft reset a group of Border Gateway Protocol (BGP) neighbors, use the **clear bgp soft** command in XR EXEC mode.

clear bgp {**ipv4** {**unicast** | **multicast** | **labeled-unicast** | **all** | **tunnel** | **mdt**} | **ipv6** {**unicast** | **multicast** | **labeled-unicast** | **all** } | **all** {**unicast** | **multicast** | **labeled-unicast** | **all** | **tunnel** | **mdt**} | **vpnv4 unicast** | **vrf** {*vrf-name* | **all**} | **vpnv6 unicast**} {***** *ip-address* | **as***as-number* | **external**}**soft**[[**in** | {**prefix-filter**} | **out**]]

Syntax Description

ipv4	Specifies IP Version 4 address prefixes.
unicast	Specifies unicast address prefixes.
multicast	Specifies multicast address prefixes.
labeled-unicast	Specifies labeled unicast address prefixes.
all	For subaddress families, specifies prefixes for all subaddress families.
tunnel	Specifies tunnel address prefixes.
mdt	Specifies IPv4 multicast distribution tree (MDT) address prefixes.
ipv6	Specifies IP Version 6 address prefixes.
all	For address family, specifies prefixes for all address families.
vpnv4 unicast	Specifies VPNv4 unicast address families.
vrf	Specifies VPN routing and forwarding (VRF).
vrf-name	Name of a VRF.
all	For VRF, specifies all VRFs.
ipv4 { **unicast | labeled-unicast** }	For VRF, specifies IPv4 unicast or labeled-unicast address families.
ipv6 unicast	For VRF, specifies IPv6 unicast address families.
vpnv6 unicast	Specifies VPNv6 unicast address families.
*****	Soft resets all BGP neighbors.
ip-address	IP address of the neighbor to be reset.
as *as-number*	Autonomous system (AS) number for all neighbors to be reset. Range for 2-byte numbers is 1 to 65535. Range for 4-byte numbers is 1.0 to 65535.65535.
external	Specifies clearing of all external peers.

in	(Optional) Triggers an inbound soft reset. If the **in** or **out** keyword is not specified, both inbound and outbound soft resets are triggered.
prefix-filter	(Optional) Specifies to send a new Outbound Route Filter (ORF) to the neighbor. Neighbor installs the new ORF and resends its routes.
out	(Optional) Triggers an outbound soft reset. If the **in** or **out** keyword is not specified, both inbound and outbound soft resets are triggered.

Command Default

No default behavior or value

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

Use the **clear bgp soft** command to trigger a soft reset of the specified address families for the specified group of neighbors. This command is useful if you change the inbound or outbound policy for the neighbors, or any other configuration that affects the sending or receiving of routing updates.

If an outbound soft reset is triggered, BGP resends all routes for the address family to the given neighbors.

If an inbound soft reset is triggered, BGP by default sends a REFRESH request to the neighbor, if the neighbor has advertised the ROUTE_REFRESH capability. To determine whether the neighbor has advertised the ROUTE_REFRESH capability, use the **show bgp neighbors** command, and look for the following line of output:

```
Received route refresh capability from peer.
```

If the neighbor does not support route refresh, but the **soft-reconfiguration inbound** command is configured for the neighbor, then BGP uses the routes cached as a result of the **soft-reconfiguration inbound** command to perform the soft reset.

If you want BGP to use the cached routes even if the neighbor supports route refresh, you can use the **always** keyword when configuring the **soft-reconfiguration inbound** command.

If the neighbor does not support route refresh and the **soft-reconfiguration inbound** command is not configured, then inbound soft reset is not possible. In this case, an error is printed.

Note By default, if the configuration for an inbound or outbound route policy is changed, BGP performs an automatic soft reset. Use the **bgp auto-policy-soft-reset disable** command to disable this behavior.

Task ID

Task ID	Operations
bgp	execute

Examples

The following example shows how to trigger an inbound soft clear for IPv4 unicast routes received from neighbor 10.0.0.1:

```
RP/0/RP0/CPU0:router# clear bgp ipv4 unicast 10.0.0.1 soft in
```

clear bgp vrf rpki validation

To clear the Border Gateway Protocol (BGP) resource public key infrastructure (RPKI) origin-validation table, use the **clear bgp rpki validation** command in XR EXEC mode.

clear bgp vrf`vrf-name` **rpki validation** [**ipv4** | **ipv6**]

Syntax Description

vrf *vrf-name*	Specifies VPN routing and forwarding (VRF) instance.
Ipv4	Clears trigger origin validation walk for the IPv4 table.
Ipv6	Clears trigger origin validation walk for the IPv6 table.

Command Default

No default behavior or values.

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 6.0	This command was introduced
Release 6.5.1	This command was modified. The **vrf** *vrf-name* keyword-argument combination were added.

Usage Guidelines

None

Task ID

Task ID	Operation
bgp	read, write

Example

TThe following example shows how to clear the BGP RPKI origin-validation table.

```
Router#clear bgp vrf vrf1 rpki validation
```

cluster-id allow-equal

To enable a router to accept BGP routes which have the same first cluster-ID as the router's own cluster-ID in the list of cluster-IDs, use the **cluster-id allow-equal** command in BGP neighbor address-family configuration mode. To revert to the default behaviour, use the **disable** option, or **no** form of this command.

cluster-id allow-equal [**disable**]

Syntax Description

disable	Specifies that BGP routes which have the same first cluster-ID as the router's own cluster-ID in the list of cluster-IDs, is rejected.

Command Default

A router rejects BGP routes which have the same first cluster-ID as the router's own cluster-ID in the list of cluster-IDs.

Command Modes

BGP neighbor address-family configuration.

Command History

Release	Modification
Release 6.6.3	This command was introduced.

Usage Guidelines

To avoid routing loops, the **cluster-id allow-equal** configuration must be done with care.

The following example shows how to enable a router to accept BGP routes which have the same first cluster-ID as the router's own cluster-ID in the list of cluster-IDs.

```
Router(config)# router bgp 100
Router(config-bgp)# neighbor 10.1.1.1
Router(config-bgp-nbr)# address-family ipv4 unicast
Router(config-bgp-nbr-af)# cluster-id allow-equal
```

The following example shows how to disable a router from accepting BGP routes from the same cluster-ID as the router's own cluster-ID in the list of cluster-IDs.

```
Router(config)# router bgp 100
Router(config-bgp)# neighbor 10.1.1.1
Router(config-bgp-nbr)# address-family ipv4 unicast
Router(config-bgp-nbr-af)# cluster-id allow-equal disable
```

default-information originate (BGP)

To allow origination of a default route to be redistributed into the Border Gateway Protocol (BGP) from another protocol, use the **default-information originate** command in an appropriate configuration mode. To disable this function, use the **no** form of this command.

default-information originate
no default-information originate

Syntax Description

This command has no arguments or keywords.

Command Default

BGP does not permit redistribution of a default route into BGP.

Command Modes

Router configuration

VRF configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

Use the **redistribute** command to redistribute routes from another protocol into BGP. By default, if these routes include the default route (0.0.0.0/0 for IPv4 or ::/0 for IPv6), the default route is ignored. Use the **default-information originate** command to change this behavior so that the default route is not ignored and is redistributed into BGP along with the other routes for the protocol being redistributed.

Task ID

Task ID	Operations
bgp	read, write

Examples

The following example shows how to configure BGP to redistribute the default route into BGP:

```
RP/0/RP0/CPU0:router(config)#router bgp 164
RP/0/RP0/CPU0:router(config-bgp)# default-information originate
```

default-martian-check disable

To disable the Martian check on the IPv4 and IPv6 prefixes, use the **default-martian-check disable** command in the address-family configuration mode. To enable the Martian check on the IPv4 and IPv6 prefixes, use the **no** form of this command.

default-martian-check disable
no default-martian-check disable

Syntax Description

This command has no keywords or arguments.

Command Default

None

Command Modes

IPv4 address family configuration mode.

IPv6 address family configuration mode.

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

Task ID

Task ID	Operations
bgp	read, write

Examples

This example shows how to disable Martian check for an IPv4 address prefix.

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router bgp 100
RP/0/RP0/CPU0:router(config-bgp)# address-family ipv4 multicast
RP/0/RP0/CPU0:router(config-bgp-af)# default-martian-check disable
```

This example shows how to disable Martian check for an IPv6 address prefix.

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router bgp 100
RP/0/RP0/CPU0:router(config-bgp)# address-family ipv6 multicast
RP/0/RP0/CPU0:router(config-bgp-af)# default-martian-check disable
```

default-metric (BGP)

To set default metric values for the Border Gateway Protocol (BGP), use the **default-metric** command in an appropriate configuration mode. To disable metric values, use the **no** form of this command.

default-metric *value*
no default-metric [*value*]

Syntax Description

value	Default metric value appropriate for the specified routing protocol. Range is 1 to 4294967295.

Command Default

A metric is not set.

Command Modes

Router configuration

VRF configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

Use the **default-metric** command to set the Multi Exit Discriminator (MED) to advertise to peers for routes that do not already have a metric set (routes that were received with no MED attribute).

Note The metric values that you apply using the **default-metric** command take effect only for a new prefix which gets into the BGP table. The metrics for the existing prefixes in the BGP table remain the same. Also, when you remove the **default-metric** command from the configuration, the metrics which were previously assigned for prefixes are not updated. To get out of this condition, clear the BGP neighborship.

Task ID

Task ID	Operations
bgp	read, write

Examples

The following example shows how to set the BGP default metric:

```
RP/0/RP0/CPU0:router(config)# router bgp 109
RP/0/RP0/CPU0:router(config-bgp)# default-metric 10
```

default-originate

To cause a Border Gateway Protocol (BGP) speaker (the local router) to send the default route 0.0.0.0/0 to a neighbor for use as a default route, use the **default-originate** command in an appropriate configuration mode. To disable this function, use the **no** form of this command.

default-originate [**inheritance-disable** | **route-policy** *route-policy-name*]
no default-originate [**inheritance-disable** | **route-policy** *route-policy-name*]

Syntax Description

inheritance-disable	(Optional) Prevents the **default-originate** command characteristics from being inherited from a parent group.
route-policy *route-policy-name*	(Optional) Specifies the name of a route policy. The route policy allows route 0.0.0.0 to be injected conditionally. IPv6 address family is supported.

Command Default

The default route is not advertised to BGP neighbors.

Command Modes

IPv4 neighbor address family configuration

IPv6 neighbor address family configuration

IPv4 neighbor group address family configuration

IPv6 neighbor group address family configuration

IPv4 address family group configuration

IPv6 address family group configuration

L2VPN EVPN address family group configuration

VRF IPv4 neighbor address family configuration

VRF IPv6 neighbor address family configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

The **default-originate** command does not require the presence of the default route (0.0.0.0/0 for IPv4 or ::/0 for IPv6) in the local router. When the **default-originate** command is used with a route policy, the default route is advertised if any route in the BGP table matches the policy.

In the L2VPN EVPN address-family group configuration mode, you can configure either **default-originate** or **default-originate inheritance-disable** commands.

Task ID

Task ID	Operations
bgp	read, write

Examples

The following example shows how to unconditionally advertise the route 0.0.0.0/0 to the neighbor 172.20.2.3:

```
RP/0/RP0/CPU0:router(config)# router bgp 109
RP/0/RP0/CPU0:router(config-bgp)#address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-bgp)# neighbor 172.20.2.3
RP/0/RP0/CPU0:router(config-bgp-nbr)# remote-as 200
RP/0/RP0/CPU0:router(config-bgp-nbr)# address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-bgp-nbr-af)# default-originate
```

The following example shows how to advertise the route 0.0.0.0/0 to the neighbor 172.20.2.3 only if a route exists in the BGP table that matches the route policy called default-default-policy:

```
RP/0/RP0/CPU0:router(config)# router bgp 109
RP/0/RP0/CPU0:router(config-bgp)# neighbor 172.20.2.3
RP/0/RP0/CPU0:router(config-bgp-nbr)# remote-as 200
RP/0/RP0/CPU0:router(config-bgp-nbr)# address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-bgp-nbr-af)# default-originate route-policy default-default-policy
```

default-prefix suppress (BGP)

To suppress automatic generation of the default RT (Route Target) membership prefix, use the **default-prefix suppress** command under the global IPv4 RT-filter address family. To restore automatic default RT membership-prefix generation when legacy trigger conditions are met, use the **no** form of this command.

default-prefix suppress

Syntax Description

default-prefix suppress	Suppresses automatic generation of the default RT membership prefix.

Command Default

Default-prefix suppression is disabled. IOS XR retains the legacy automatic generation behavior when the applicable trigger conditions are met.

Command Modes

BGP IPv4 RT-filter address family configuration mode

Command History

Release	Modification
Release 26.3.1	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

Task ID

Task ID	Operations
bgp	read, write

Examples

This example shows how to enable default RT prefix suppression.

```
Router# config
Router(config)# router bgp 65
Router(config-bgp)# address-family ipv4 rt-filter
Router(config-bgp-af)# default-prefix suppress
Router(config-bgp-af)# commit
```

description (BGP)

To annotate a neighbor, neighbor group, VPN routing and forwarding (VRF) neighbor, or session group, use the **description** command in an appropriate configuration mode. To remove the annotation, use the **no** form of this command.

description *text*
no description [*text*]

Syntax Description

text	Meaningful description or comment. Maximum of 80 characters.

Command Default

No comment or description exists.

Command Modes

Neighbor group configuration

Neighbor configuration

Session group configuration

VRF neighbor configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

Use the **description** command to provide a description of a neighbor, neighbor group, VRF neighbor, or session group. The description is used to save user comments and does not affect software function.

Task ID

Task ID	Operations
bgp	read, write

Examples

The following example shows how to configure the description “Our best customer” on the neighbor 192.168.13.4:

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)#router bgp 65000
RP/0/RP0/CPU0:router(config-bgp)#neighbor 192.168.13.4
RP/0/RP0/CPU0:router(config-bgp-nbr)#description Our best customer
```

distance bgp

To allow the use of external, internal, and local administrative distances that could be used to prefer one class of routes over another, use the **distance bgp** command in an appropriate configuration mode. To disable the use of administrative distances, use the **nono** form of this command.

distance bgp *external-distance internal-distance local-distance*
no distance bgp [*external-distance internal-distance local-distance*]

Syntax Description

external-distance	Administrative distance for Border Gateway Protocol (BGP) external routes. External routes are routes for which the best path is learned from a neighbor external to the autonomous system. Range is 1 to 255. Routes with a distance of 255 are not installed in the routing table.
internal-distance	Administrative distance for BGP internal routes. Internal routes are those routes that are learned from another BGP entity within the same autonomous system. Range is 1 to 255. Routes with a distance of 255 are not installed in the routing table.
local-distance	Administrative distance for BGP local routes. The *local-distance* argument applies to locally generated aggregate routes (such as the routes generated by the **aggregate-address** command) and backdoor routes installed in the routing table. Range is 1 to 255. Routes with a distance of 255 are not installed in the routing table.

Command Default

external-distance : 20

internal-distance : 200

local-distance : 200

Command Modes

IPv4 address family configuration

IPv6 address family configuration

VRF IPv4 address family configuration

VRF IPv6 address family configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

Use the **distance bgp** command if another protocol is known to be able to provide a better route to a node than was actually learned using external BGP, or if some internal routes should be preferred by BGP.



Note Changing the administrative distance of BGP internal routes is considered risky and is not recommended. One problem that can arise is the accumulation of routing table inconsistencies, which can interfere with routing.

An administrative distance is a rating of the trustworthiness of a routing information source. Numerically, an administrative distance is an integer from 1 to 255. In general, the higher the value, the lower the trust rating. An administrative distance of 255 means the routing information source cannot be trusted at all and should be ignored.

Task ID

Task ID	Operations
bgp	read, write

Examples

The following example shows that iBGP routes are preferable to locally generated routes, so the administrative distance values are set accordingly:

```
RP/0/RP0/CPU0:router(config)# router bgp 109
RP/0/RP0/CPU0:router(config-bgp)# address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-bgp-af)#distance bgp 20 20 200
```

distribute bgp-ls (ISIS)

To distribute ISIS link-state data using BGP LS, use the **distribute bgp-ls** command in XR Config mode. To stop link-state distribution, use the **no** form of this command.

distribute bgp-ls [instance-id *value***] [level** {**1** | **2**}**] [throttle** *time***]**
no distribute bgp-ls

Syntax Description

instance-id *value*	(Optional) Specifies the instance identifier defined by the **router isis** command. Range is from 1 to 65535. If the instance-id is not configured, the system assigned instance-id for the ISIS process will be used.
level 1 | 2	(Optional) Displays IS-IS link-state database for Level 1 or Level 2 independently.
throttle	(Optional) Specifies throttle update, in seconds. Range is from 5 to 20 seconds.

Command Default

None

Command Modes

XR Config mode

Command History

Release	Modification
Release 6.1.1	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

Task ID

Task ID	Operations
isis	read, write

Examples

This example shows how to distribute ISIS link-state information using BGP LS:

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router isis foo
RP/0/RP0/CPU0:router(config-isis)# distribute bgp-ls instance-id 32 level 2 throttle 5
```

distribute bgp-ls (OSPF)

To distribute OSPFv2 and OSPFv3 link-state data using BGP LS, use the **distribute bgp-ls** command in XR Config mode. To stop link-state distribution, use the **no** form of this command.

distribute bgp-ls [instance-id *value***] [throttle** *time***]**
no distribute bgp-ls

Syntax Description

instance-id *value*	(Optional) Specifies the instance identifier defined by the **router ospf** command. Range is from 1 to 65535. If the instance-id is not configured, the system assigned instance-id for the OSPF process is used.
throttle	(Optional) Specifies throttle time between successive link-state advertisement (LSA) updates. Range is from 0 to 3600.

Command Default

BGP distribution is disabled.

Command Modes

XR Config mode

Command History

Release	Modification
Release 6.1.1	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

Task ID

Task ID	Operations
ospf	read, write

Examples

This example shows how to distribute OSPF link-state information using BGP LS:

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router ospf 100
RP/0/RP0/CPU0:router(config-ospf)# distribute bgp-ls instance-id 32 throttle 10
```

domain-distinguisher

To configure globally unique identifier ASN for IGP domain, use the **domain-distinguisher** command in address-family link-state configuration mode. To remove unique identifier, use the **no** form of this command.

domain-distinguisher *unique-id*
no domain-distinguisher

Syntax Description

unique-id	Specifies four-octet unique identifier ASN. Range is from 1 to 4294967295.

Command Default

None

Command Modes

Address-family link-state configuration.

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

Task ID

Task ID	Operations
bgp	read, write

Examples

This example shows how to configure a unique identifier ASN:

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router bgp 100
RP/0/RP0/CPU0:router(config-bgp)# address-family link-state link-state
RP/0/RP0/CPU0:router(config-bgp-af)# domain-distinguisher 1234
```

dmz-link-bandwidth

To originate a demilitarized zone (DMZ) link bandwidth extended community for the link to an eBGP neighbor, use the **dmz-link-bandwidth** command in an Neighbor configuration mode. To stop origination of the DMZ link bandwidth extended community, use the **no** form of this command.

dmz-link-bandwidth [**inheritance-disable**]
no dmz-link-bandwidth

Syntax Description

inheritance-disable	(Optional) Prevents the **dmz-link-bandwidth** command from being inherited from a parent group.

Command Default

BGP does not originate the DMZ link bandwidth extended community.

Command Modes

Neighbor configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

Use the **dmz-link-bandwidth** command to advertise the bandwidth of links that are used to exit an autonomous system.

Task ID

Task ID	Operations
bgp	read, write

Examples

This example shows how to advertise the bandwidth of links to eBGP neighbors from router bgp 1:

```
RP/0/RP0/CPU0:router(config)# router bgp 1
RP/0/RP0/CPU0:router(config-bgp)#neighbor 10.67.89.01
RP/0/RP0/CPU0:router(config-bgp-nbr)#dmz-link-bandwidth
```

dscp (BGP)

To set the differentiated services code point (DSCP) value, use the **dscp** command in the appropriate configuration mode. To remove the **dscp** command from the configuration file and restore the system to its default interval values, use the no form of this command.

dscp *value*
no dscp [*value*]

Syntax Description

value	Value of the DSCP. The DSCP value can be a number from 0 to 63, or it can be one of the following keywords: **default** , **ef** , **af11** , **af12** , **af13** , **af21** , **af22** , **af23** , **af31** , **af32** , **af33** , **af41** , **af42** , **af43** , **cs1** , **cs2** , **cs3** , **cs4** , **cs5** , **cs6** , or **cs7** .

Command Default

No default behavior or values

Command Modes

Neighbor configuration

Neighbor session group configuration

Neighbor group configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

Use the **dscp** command to change the minimum and maximum packet thresholds for the DSCP value.

Table 2: dscp Default Settings, on page 152 lists the DSCP default settings used by the **dscp** command. The DSCP value, corresponding minimum threshold, maximum threshold, and mark probability are listed. The last row of the table (the row labeled "default") shows the default settings used for any DSCP value not specifically shown in the table.

Table 2: dscp Default Settings

DSCP (Precedence)	Minimum Threshold	Maximum Threshold	Mark Probability
af11	32	40	1/10
af12	28	40	1/10
af13	24	40	1/10
af21	32	40	1/10
af22	28	40	1/10
af23	24	40	1/10
af31	32	40	1/10

DSCP (Precedence)	Minimum Threshold	Maximum Threshold	Mark Probability
af32	28	40	1/10
af33	24	40	1/10
af41	32	40	1/10
af42	28	40	1/10
af43	24	40	1/10
cs1	22	40	1/10
cs1	24	40	1/10
cs3	26	40	1/10
cs4	28	40	1/10
cs5	30	40	1/10
cs6	32	40	1/10
cs7	34	40	1/10
ef	36	40	1/10
default	20	40	1/10

Task ID

Task ID	Operations
bgp	read, write

Examples

The following example shows how to set the DSCP value to af32:

```
RP/0/RP0/CPU0:router(config)# router bgp 5
RP/0/RP0/CPU0:router(config-bgp)#neighbor 10.1.1.1
RP/0/RP0/CPU0:router(config-bgp-nbr)#remote-as 100
RP/0/RP0/CPU0:router(config-bgp-nbr)# dscp af32
```

dynamic-med interval

To customize advertisement of Route Distinguisher (RD) updates to neighbors, use the **dynamic-med-interval** command in BGP address family configuration mode for IPv4, IPv6, VPNv4, and VPNv6 unicast address families. To undo this command configuration, use the **no** form of this command.

dynamic-med interval *interval-value*
no dynamic-med interval *interval-value*

Syntax Description

interval *interval-value*	Specifies interval (in seconds) between two consecutive update generation to neighbors.

Command Default

The default value of dynamic-med interval is 600 seconds.

Command Modes

BGP address family configuration mode

Command History

Release	Modification
Release 5.3.2	This command was introduced.

Task ID

Task ID	Operation
bgp	read, write

Example

The following example shows how to customize rd-advertisement of updates to neighbors:

```
Router# configure
Router(config)# router bgp 100
Router(config-bgp)# address-family ipv4 unicast
Router(config-bgp-af)# dynamic-med interval 2
Router(config-bgp-nbr-af)# commit
```

ebgp-multihop

To accept and attempt Border Gateway Protocol (BGP) connections to external peers residing on networks that are not directly connected, use the **ebgp-multihop** command in an appropriate configuration mode. To disable connections to external peers and allow only direct connections between neighbors, use the **no** form of this command.

ebgp-multihop [*ttl-value*] [**mpls**]
no ebgp-multihop [*ttl-value*] [**mpls**]

Syntax Description

ttl-value	(Optional) Time-to-live (TTL) value. Range is 1 to 255 hops.
mpls	(Optional) Disables BGP label rewrite.

Command Default

Default TTL value is 255.

Command Modes

Neighbor configuration

VRF neighbor configuration

Neighbor group configuration

Session group configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

Use the **ebgp-multihop** command to enable multihop peerings with external BGP neighbors. The BGP protocol states that external neighbors must be directly connected (one hop away). The software enforces this by default; however, the **ebgp-multihop** command can be used to override this behavior.

Use of the **mpls** option in the **ebgp-multihop** command prevents BGP from enabling MPLS on the peering interface and also prevents allocation of Implicit-NULL rewrite labels for nexthop addresses learned from the peer. This is useful in some scenarios in which MPLS forwarding labels to the nexthops have already been learned via BGP labeled-unicast or LDP.

If this command is configured for a neighbor group or session group, all neighbors using the group inherit the configuration. Values of commands configured specifically for a neighbor override inherited values.

Task ID

Task ID	Operations
bgp	read, write

Examples

The following example shows how to allow a BGP connection to neighbor 172.20.16.6 of up to 255 hops away:

```
RP/0/RP0/CPU0:router(config)# router bgp 109
RP/0/RP0/CPU0:router(config-bgp)# neighbor 172.20.16.6
RP/0/RP0/CPU0:router(config-bgp-nbr)# ebgp-multihop
```

export route-policy

To configure an export route policy, use the **export route-policy** command in an appropriate configuration mode. To restore the system to its default condition, use the **no** form of this command.

export route-policy *policy-name*
no export route-policy [*policy-name*]

Syntax Description

policy-name	Name of the configured route policy.

Command Default

No default behavior or values

Command Modes

Global VRF IPv4 address family configuration

Global VRF IPv6 address family configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

Use the **export route-policy** command to define the conditions that allow specified routes to be tagged with specified route-targets.

You cannot use this **route-policy** to drop a route.

Task ID

Task ID	Operations
bgp	read, write
ip-services	read, write

Examples

The following example shows how to configure an export route policy:

```
RP/0/RP0/CPU0:router(config)# vrf vrf-1
RP/0/RP0/CPU0:router(config-vrf)#address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-vrf-af)# export route-policy policy-A
```

export route-target

To configure a VPN routing and forwarding (VRF) export route-target extended community, use the **export route-target** command in an appropriate configuration mode. To restore the system to its default condition, use the **no** form of this command.

export route-target [*as-number:nn* *ip-address:nn*]
no export route-target [*as-number:nn* *ip-address:nn*]

Syntax Description

as-number:nn	(Optional) *as-number* —Autonomous system (AS) number of the route-target extended community. • as-number • Range for 2-byte Autonomous system numbers (ASNs) is 1 to 65535. • Range for 4-byte Autonomous system numbers (ASNs) in asplain format is 1 to 4294967295. • Range for 4-byte Autonomous system numbers (ASNs) is asdot format is 1.0 to 65535.65535. • *nn* —32-bit number
ip-address:nn	(Optional) IP address of the route-target extended community. • *ip-address* —32-bit IP address • *nn* —16-bit number

Command Default

No default behavior or values

Command Modes

Global VRF IPv4 address family configuration

Global VRF IPv6 address family configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

Export route-target extended communities are associated with prefixes when advertised to remote provider edge (PE) routers. The remote PE routers import the route-target extended communities into a VRF instance that has the import route-targets that match the exported route-target extended communities.

To specify multiple route targets, enter export route target configuration mode then enter one route target for each command line.

Task ID

Task ID	Operations
bgp	read, write

Task ID	Operations
ip-services	read, write

Examples

The following example shows how to specify an export route-target:

```
RP/0/RP0/CPU0:router(config)# vrf vrf-1
RP/0/RP0/CPU0:router(config-vrf)# address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-vrf-af)# export route-target 500:1
```

hw-module fib bgp-mp-pic-auto-protect-enable

To enable auto protection feature for BGP-LU multipath PIC, use the **hw-module fib bgp-mp-pic auto-protect***enable* command in XR Config Mode.

hw-module fib bgp-mp-pic auto-protect*enable*

Syntax Description

enable	To enable auto-protection feature.

Command Default

Auto-protection feature for BGP-LU multipath is disabled by default.

Command Modes

XR Config Mode

Command History

Release	Modification
Release 7.5.4, Release 7.8.2, Release 7.9.1	This command was introduced.

Usage Guidelines

After executing the **hw-module fib bgp-mp-pic auto-protect enable** command, you must reload the router.

Task ID

Task ID	Operation
bgp	read, write

The following example shows how to enable the auto-protection feature for BGP-LU multipath PIC.

```
RP/0/RP0/CPU0:ios#configure
Wed Mar  8 13:56:48.693 UTC
RP/0/RP0/CPU0:ios(config)#hw-module fib bgp-mp-pic auto-protect enable
Wed Mar  8 13:56:54.449 UTC
In order to activate/deactivate bgp multipath auto protect, you must manually reload the
chassis/all line cards
RP/0/RP0/CPU0:ios(config)#commit
Wed Mar  8 13:56:55.348 UTC
RP/0/RP0/CPU0:ios(config)#end
RP/0/RP0/CPU0:ios#
```

hw-module fib dlb level-1 enable

To enable the router to support a higher scale (more than 2K) of labelled prefixes with Equal Cost Multi Path (ECMP), use the command **hw-module fib dlb level-1 enable**.

hw-module fib dlb level-1 enable

Syntax Description

This command has no keywords or arguments.

Command Default

None.

Command Modes

XR Config mode

Command History

Release	Modification
Release 7.3.1	This command was introduced.

Usage Guidelines

- Make sure that you reload the line card for this configuration to take effect.
- This table provides the list of supported use cases.

Table 3: Supported Use Cases

Use Case	Support
BGP L3VPN over IGP DLB	YES
BGP LU over IGP DLB	YES
BGP PIC edge convergence with IGP+DLB	NO
IGP DLB over SRTE	NO
IGP DLB with LDP	YES
IGP DLB with SR	YES
L2VPN PW over IGP DLB	YES
TI-LFA/LFA/RLFA at IGP with DLB	YES
Other features (not listed above)	NO

Task ID

Task ID	Operation
config-services	read, write
root-lr	read, write

Example

The following example shows how to enable the router to support a higher scale (> 2K) of labelled prefixes with ECMP.

```
Router# configure
Router(config)# hw-module fib dlb level-1 enable
Router(config-vrf)# commit
Router(config)# end
Router# reload location all
Proceed with reload? [confirm]
```

hw-module profile stats j2-dynamic-stats

To enhance the number of flows to 16K, use the **hw-module profile stats j2-dynamic-stats** command in XR Config mode.

hw-module profile stats j2-dynamic-stats

Syntax Description

This command has no keywords or arguments.

Command Default

None

Command Modes

XR Config

Command History

Release	Modification
Release 7.6.1	This command was deprecated.
Release 7.4.1	This command was introduced.

Usage Guidelines

The router must be reloaded for the configuration of the **hw-module profile stats j2-dynamic-stats** command to take effect.

Example

The following example shows how to enhance the number of flows to 16K:

```
Router(config)# hw-module profile stats j2-dynamic-stats
```

hw-module fib bgppa stats-mode

To enable the BGP policy accounting on the main interface or on the sub interface, run the **hw-module fib bgppa stats-mode** command with the **main-intf** or the **sub-intf** keywords respectively.

hw-module fib bgppa stats-mode {main-intf | sub-intf}

Syntax Description

bgppa	BGP policy accounting
stats-mode	Statistics accounting mode
main-intf	BGP policy accounting statistics for the main interface.
sub-intf	BGP policy accounting statistics for the sub interface.

Command Default

None

Command Modes

Interface configuration

Command History

Release	Modification
Release 7.9.1	This command was introduced.

Usage Guidelines

- The BGP policy accounting feature is applicable for the following address families:
  - IPv4
  - IPv6
- After configuring the command, you must reload the router for the BGP policy accounting feature to take effect.

Task ID

Task ID	Operation
bgp	read, write

The following example shows the configuration of BGP policy accounting on the main interface and the sub interface. You must reload the router after configuring the following commands to take effect.

For main interface:

```
Router# config
Router(config)# hw-module fib bgppa stats-mode main-intf
In order to enable BGPPA only in J2 cards with external TCAM, you must manually reload the
 chassis/all line cards.
Router(config)# commit
```

For sub interface:

```
Router# config
Router(config)# hw-module fib bgppa stats-mode sub-intf
In order to enable BGPPA only in J2 cards with external TCAM, you must manually reload the
 chassis/all line cards.
Router(config)# commit
```

hw-module fib urpf disable

To disable the Unicast Reverse Path Forwarding (URPF) infrastructure globally on Cisco NCS 5500 series modular routers having Cisco NCS 5700 line cards, use the **hw-module fib urpf disable** command in the Configuration mode.

hw-module fib urpf disable

Syntax Description

This command has no keywords or arguments.

Command Default

None

Command Modes

Global Configuration mode

Command History

Release	Modification
Release 7.5.1	This command was introduced.

Usage Guidelines

After executing the **hw-module fib urpf disable** command, you must reload the router.

The following example shows how to disable the urpf infrastructure globally in a NCS 5700 line card:

```
Router#config
Router(config)#hw-module fib urpf disable
Router(config)#commit
```

ibgp policy out enforce-modifications

To allow an outbound route policy for an internal BGP (iBGP) peer to modify all BGP route attributes, only when an iBGP route is sent to another iBGP peer (only on route-reflectors), use the **ibgp policy out enforce-modifications** command in XR Config mode. To disable this feature, use the **no** form of this command.

ibgp policy out enforce-modifications
no ibgp policy out enforce-modifications

Syntax Description

This command has no arguments or keywords.

Command Default

ibgp policy out enforce-modifications is disabled.

Command Modes

XR Config mode

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

Use the **ibgp policy out enforce-modifications** command to set and modify BGP route attributes for updates to iBGP peers.

If the **ibgp policy out enforce-modifications command** is configured under router BGP configuration, then all the changes made by the outbound policy for an iBGP peer will be present in an update message sent to the peer.

Task ID

Task ID	Operations
bgp	read, write

Examples

The following example shows how to set the ibgp policy out enforce-modifications:

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router bgp 6500
RP/0/RP0/CPU0:router(config-bgp)# ibgp policy out enforce-modifications
```

import route-policy

To configure an import route policy, use the **import route-policy** command in an appropriate configuration mode. To restore the system to its default condition, use the **no** form of this command.

import route-policy *policy-name*
no import route-policy [*policy-name*]

Syntax Description

policy-name	Name of the configured route policy.

Command Default

No default behavior or values

Command Modes

Global VRF IPv4 address family configuration

Global VRF IPv6 address family configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

Use the **import route-policy** command to define the conditions that allow specified routes to be imported into the VPN routing and forwarding (VRF) instance if the routes are tagged with specified route-targets.

Task ID

Task ID	Operations
bgp	read, write
ip-services	read, write

Examples

The following example shows how to allow only policy-B to be imported to VRF:

```
RP/0/RP0/CPU0:router(config)# vrf vrf-1
RP/0/RP0/CPU0:router(config-vrf)# address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-vrf-af)# import route-policy policy-B
```

import route-target

To configure a VPN routing and forwarding (VRF) import route-target extended community, use the **import route-target** command in an appropriate configuration mode. To restore the system to its default condition, use the **no** form of this command.

import route-target [*as-number:nn ip-address:nn*]
noimport route-target [*as-number:nn ip-address:nn*]

Syntax Description

as-number:nn	(Optional) Autonomous system (AS) number of the route-target extended community. • *as-number* • Range for 2-byte Autonomous system numbers (ASNs) is 1 to 65535. • Range for 4-byte Autonomous system numbers (ASNs) in asplain format is 1 to 4294967295. • Range for 4-byte Autonomous system numbers (ASNs) is asdot format is 1.0 to 65535.65535. • *nn* —32-bit number
ip-address:nn	(Optional) IP address of the route-target extended community. • *ip-address* —32-bit IP address • *nn* —16-bit number

Command Default

No default behavior or values

Command Modes

Global VRF IPv4 address family configuration

Global VRF IPv6 address family configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

Use the **import route-target** command to specify that prefixes associated with the configured import route-target extended communities are imported into the VRF instance.

To specify multiple route targets, enter import route target configuration mode, then enter one route target for each command line.

Task ID

Task ID	Operations
bgp	read, write
ip-services	read, write

Examples

The following example shows how to specify an import route-target:

```
RP/0/RP0/CPU0:router(config)#vrf vrf-1
RP/0/RP0/CPU0:router(config-vrf)# address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-vrf-af)# import route-target 500:99
```

ignore-connected-check

To enable the software to bypass the directly connected next hop check for single-hop eBGP peering, use the **ignore-connected-check** command in an appropriate configuration mode. To re-enable the directly connected next hop check, use the **no** form of this command.

ignore-connected-check [**inheritance-disable**]
no ignore-connected-check

Syntax Description

inheritance-disable	Prevents the **ignore-connected-check** command from being inherited from the parent.

Command Default

Ability to bypass the directly connected next hop check is disabled.

Command Modes

Neighbor configuration

Neighbor group configuration

Session group configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

Task ID

Task ID	Operations
bgp	read, write

Examples

The following example shows how to enable ignore-connected check configuration for neighbor 10.2.3.4:

```
RP/0/RP0/CPU0:router(config)# router bgp 100
RP/0/RP0/CPU0:router(config-bgp)# neighbor 10.2.3.4
RP/0/RP0/CPU0:router(config-bgp-nbr)# ignore-connected-check
```

graceful-maintenance

To allow the network to perform convergence before the router or link is taken out of service, use the **graceful-maintenance** command in the router BGP, neighbor or neighbor group configuration mode, as appropriate. To disable the command, use the **no** form of this command.

graceful-maintenance activate [**all-neighbors** | **retain-routes**]

Note This command is executed in the router BGP configuration mode.

graceful-maintenance {**activate** [**as-prepends** *as-prepends-value*] [**inheritance-disable**] | [**local-preference** *local-pref-value*] **inheritance-disable**}

Note This command is executed in either the neighbor configuration or neighbor group configuration mode.

Syntax Description

activate	Announces routes with the graceful maintenance attributes while activated either under the neighbor or router BGP configuration. While activated, all routes to this neighbor are announced with the attribute configured here and all routes from this neighbor are announced to other neighbors with the graceful maintenance attributes configured under those neighbors. The GSHUT community is announced regardless of the other attributes configured here. To allow the GSHUT community to be announced to eBGP neighbors, you must configure the **send-community-gshut-ebgp** command.
all-neighbors	If you use the **all-neighbors** keyword, Graceful Maintenance is activated even for those neighbors that do not have Graceful Maintenance activated.
retain-routes	Choosing **retain-routes** causes RIB to retain BGP routes when the BGP process is stopped. You would use **retain-routes** when only BGP is being brought down instead of the entire router and if it is known that neighboring routers are being kept in operation during the maintenance of the local BGP. If RIB has alternative routes provided by another protocol or a default route, then it is recommended not to retain BGP routes after the BGP process stops.

as-prepends *as-prepends-value* **inheritance-disable**	Indicates the number of times to prepend the local AS number to the AS path of routes. The default value is 0. The keyword **inheritance-disable** prevents AS prepends from being inherited from the parent. Specifies the number of times to prepend the local AS number to the AS path of routes and advertises the GSHUT community with the local preference value specified for the routes. When the router adds the GSHUT community to a route as it advertises it, it also changes the LOCAL_PREF attribute and prepends the local AS number as specified in the commands. Sending GSHUT provides flexibility in how neighboring routers handle the lower preference: they can match it in a route policy and do the most appropriate thing with it. On the other hand, in simple networks, it is recommended to set local-preference to 0, rather than to create route policies everywhere else. **Note** LOCAL_PREF is not sent to real eBGP neighbors, but sent to confederation member AS eBGP neighbors. To lower preference to eBGP neighbors, as-prepends is required.
local-preference *local-pref-value* **inheritance-disable**	Indicates the range of values for Local Preference. The keyword **inheritance-disable** prevents local preference from being inherited from the parent.

Command Default

None

Command Modes

router BGP

neighbor configuration

neighbor group configuration

Command History

Release	Modification
Release 6.5.2	This command was introduced.

Task ID

Task ID	Operations
bgp	read, write

hw-module profile flowspec

To enable flowspec processing with IPv6 traffic, use the **hw-module profile flowspec** command in the global configuration mode. To disable flowspec, use the **no** form of this command.

hw-module profile flowspec v6-enable [**location** *node-id*]

no hw-module profile flowspec v6-enable [**location** *node-id*]

Syntax Description

v6-enable	Enables the flowspec with IPv6 traffic.
location *node-id*	(Optional) Specifies the flowspec location. The *node-id* argument is entered in the `rack/slot/module` notation

Command Default

None

Command Modes

Global configuration

Command History

Release	Modification
Release 6.5.1	This command was introduced.

Usage Guidelines

- When you enable the **hw-module profile flowspec v6-enable** command, the packets per second (PPS) rate reduces. This reduction in PPS causes both IPv6 and IPv4 line rate degradation from 835Mpps to ~700Mpps.
- This command is applicable only for the Cisco NCS 5500 Series Router.
- After executing the **hw-module profile flowspec v6-enable** command, you must reload the router.

Task ID

Task ID	Operations
bgp	read, write

Examples

The following example shows how to enable flowspec processing with IPv6 traffic:

```
Router# configure
Router# hw-module profile flowspec v6-enable
```

hw-module fib bgppa stats-mode

To enable the BGP policy accounting on the main interface or on the sub interface, run the **hw-module fib bgppa stats-mode** command with the **main-intf** or the **sub-intf** keywords respectively.

hw-module fib bgppa stats-mode {main-intf | sub-intf}

Syntax Description

bgppa	BGP policy accounting
stats-mode	Statistics accounting mode
main-intf	BGP policy accounting statistics for the main interface.
sub-intf	BGP policy accounting statistics for the sub interface.

Command Default

None

Command Modes

Interface configuration

Command History

Release	Modification
Release 7.9.1	This command was introduced.

Usage Guidelines

- The BGP policy accounting feature is applicable for the following address families:
  - IPv4
  - IPv6
- After configuring the command, you must reload the router for the BGP policy accounting feature to take effect.

Task ID

Task ID	Operation
bgp	read, write

The following example shows the configuration of BGP policy accounting on the main interface and the sub interface. You must reload the router after configuring the following commands to take effect.

For main interface:

```
Router# config
Router(config)# hw-module fib bgppa stats-mode main-intf
In order to enable BGPPA only in J2 cards with external TCAM, you must manually reload the
 chassis/all line cards.
Router(config)# commit
```

For sub interface:

```
Router# config
Router(config)# hw-module fib bgppa stats-mode sub-intf
In order to enable BGPPA only in J2 cards with external TCAM, you must manually reload the
 chassis/all line cards.
Router(config)# commit
```

keychain (BGP)

To apply key chain-based authentication on a TCP connection between two Border Gateway Protocol (BGP) neighbors, use the **keychain** command in an appropriate configuration mode. To disable key chain authentication, use the **no** form of this command.

keychain *name*
no keychain [*name*]

Syntax Description

name	Key chain name configured using the **keychain** command. The name must be a maximum of 32 alphanumeric characters.

Command Default

When this command is not specified in the appropriate configuration mode, key chain authentication is not enabled on a TCP connection between two BGP neighbors.

Command Modes

Neighbor configuration

Neighbor group configuration

Session group configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

Specify a key chain to enable key chain authentication between two BGP peers. Use the **keychain** command to implement hitless key rollover for authentication.

If this command is configured for a neighbor group or a session group, a neighbor using the group inherits the configuration. Values of commands configured specifically for a neighbor override inherited values.

Note BGP only supports HMAC-MD5 and HMAC-SHA1-12 cryptographic algorithms.

Task ID

Task ID	Operations
bgp	read, write

Examples

The following example shows how to configure neighbor 172.20.1.1 to use the key chain authentication configured in the keychain_A key chain:

```
RP/0/RP0/CPU0:router(config)# router bgp 140
RP/0/RP0/CPU0:router(config-bgp)# neighbor 172.20.1.1
RP/0/RP0/CPU0:router(config-bgp-nbr)# remote-as 1
```

```
RP/0/RP0/CPU0:router(config-bgp-nbr)# keychain keychain_A
```

keychain-disable

Note Effective with Release 3.9.0, the **keychain-disable** command was replaced by the **keychain inheritance-disable** command. See the keychain inheritance-disable, on page 181 command for more information.

To override any inherited key chain configuration from a neighbor group or session group for Border Gateway Protocol (BGP) neighbors, use the **keychain-disable** command in an appropriate configuration mode. To disable overriding any inherited key chain command, use the **no** form of this command.

keychain-disable
no keychain-disable

Syntax Description

This command has no arguments or keywords.

Command Default

Configured key chains for neighbor and session groups are inherited.

Command Modes

Neighbor configuration

Neighbor group configuration

Session group configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

If you specify a key chain on a neighbor group or session group, all users of the group inherit the key chain. Specifying a different **keychain** command specifically on a neighbor that uses the group overrides the inherited value. Specifying **keychain-disable** on a neighbor that uses the group disables key chain authentication for the neighbor.

Task ID

Task ID	Operations
bgp	read, write

Examples

The following example shows how to disable key chain authentication for neighbor 172.20.1.1, preventing it from inheriting the key chain keychain_A from session group group1:

```
RP/0/RP0/CPU0:router(config)# router bgp 140
RP/0/RP0/CPU0:router(config-bgp)# session-group group1
RP/0/RP0/CPU0:router(config-bgp-sngrp)# keychain keychain_A
RP/0/RP0/CPU0:router(config-bgp-sngrp)# exit
RP/0/RP0/CPU0:router(config-bgp)#neighbor 172.20.1.1
RP/0/RP0/CPU0:router(config-bgp-nbr)# remote-as 2
```

```
RP/0/RP0/CPU0:router(config-bgp-nbr)#use session-group group1
RP/0/RP0/CPU0:router(config-bgp-nbr)# keychain-disable
```

keychain inheritance-disable

To override any inherited key chain configuration from a neighbor group or session group for Border Gateway Protocol (BGP) neighbors, use the **keychain inheritance-disable** command in an appropriate configuration mode. To disable overriding any inherited key chain command, use the **no** form of this command.

keychain inheritance-disable
no keychain inheritance-disable

Syntax Description

This command has no arguments or keywords.

Command Default

Configured key chains for neighbor and session groups are inherited.

Command Modes

Neighbor configuration

Neighbor group configuration

Session group configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

If you specify a key chain on a neighbor group or session group, all users of the group inherit the key chain. Specifying a different **keychain** command specifically on a neighbor that uses the group overrides the inherited value. Specifying **keychain inheritance-disable** on a neighbor that uses the group disables key chain authentication for the neighbor.

Task ID

Task ID	Operations
bgp	read, write

Examples

The following example shows how to disable key chain authentication for neighbor 172.20.1.1, preventing it from inheriting the key chain keychain_A from session group group1:

```
RP/0/RP0/CPU0:router(config)#router bgp 140
RP/0/RP0/CPU0:router(config-bgp)# session-group group1
RP/0/RP0/CPU0:router(config-bgp-sngrp)# keychain keychain_A
RP/0/RP0/CPU0:router(config-bgp-sngrp)# exit
RP/0/RP0/CPU0:router(config-bgp)# neighbor 172.20.1.1
RP/0/RP0/CPU0:router(config-bgp-nbr)# remote-as 2
RP/0/RP0/CPU0:router(config-bgp-nbr)# use session-group group1
RP/0/RP0/CPU0:router(config-bgp-nbr)# keychain inheritance-disable
```

label-allocation-mode

To set the MPLS/VPN label allocation mode, use the **label-allocation-mode** command in VRF configuration mode. To remove the **label-allocation-mode** command from the configuration file and restore the system to its default condition, use the **no** form of this command.

label-allocation-mode [**per-ce** | **per-vrf**]
no label-allocation-mode

Syntax Description

per-ce	Specifies that the same label is used for all the routes advertised from a unique customer edge (CE) peer or router.
per-vrf	Specifies that the same label is used for all the routes advertised from a unique VRF.

Command Default

Per-prefix is the default label allocation mode.

Command Modes

VRF configuration

Command History

Release	Modification
Release 7.4.1	This command was deprecated
Release 6.0	This command was introduced.

Note The function of this deprecated command can be carried out using label mode command under configured address-family.

Usage Guidelines

Each prefix that belongs to a VRF instance is advertised with a single label, causing an additional lookup to be performed in the VRF forwarding table to determine the customer edge (CE) next hop for the packet. Use the **label-allocation-mode** command with the **per-ce** keyword to avoid the additional lookup on the PE router and conserve label space. This mode allows the PE router to allocate one label for every immediate next hop. The label is directly mapped to the next hop so there is no VRF route lookup performed during data forwarding. However, the number of labels allocated is one for each CE rather than one for each prefix.

Note

- The **label-allocation-mode** under the global IPv6 address family configuration mode is renamed as **label mode**, in Cisco IOS-XR Software release 4.3.1 and later releases.
- With the introduction of **label mode** command, the nexthop labels will no longer be released, when **label-allocation-mode** command with the **per-ce** keyword is unconfigured.

Task ID

Task ID	Operations
bgp	read, write

Examples

The following example shows how to set the label allocation mode to customer edge:

```
RP/0/RP0/CPU0:router(config)# router bgp 109
RP/0/RP0/CPU0:router(config-bgp)# vrf vrf-1
RP/0/RP0/CPU0:router(config-bgp-vrf)# label-allocation-mode per-ce
```

The following example shows how to set the label allocation mode to VRF:

```
RP/0/RP0/CPU0:router(config)# router bgp 109
RP/0/RP0/CPU0:router(config-bgp)# vrf vrf-1
RP/0/RP0/CPU0:router(config-bgp-vrf)# label-allocation-mode per-vrf
```

label mode

To set the MPLS/VPN label mode based on prefix value, use the **label mode** command in an appropriate configuration mode. To remove the **label mode** command from the configuration file and restore the system to its default condition, use the **no** form of this command.

Use this syntax for VPNv4 and VPNv6 unicast address-family configuration mode:

label mode per-nexthop-received-label [allocate-secondary-label]

Use this syntax for **vrf all** configuration mode under VPN IPv4/IPv6 AF (address family) mode or global IPv6 AF configuration mode:

label mode { **per-ce** | **per-vrf** | **route-policy** }
no label mode { **per-ce** | **per-vrf** | **route-policy** }

Use this syntax for IPv4/IPv6 AF configuration mode under vrf mode:

label mode { **per-prefix** | **per-ce** | **per-vrf** | **route-policy** }
no label mode { **per-prefix** | **per-ce** | **per-vrf** | **route-policy** }

Syntax Description

per-nexthop-received-label	Specifies that the same label is used for all prefixes received with the same next-hop and same received-label.
allocate-secondary-label	Enables allocating a secondary label in addition to the local label at the route reflectors to avoid label churn issue in Prefix Independent Convergence (PIC) scenario.
per-ce	Specifies that the same label is used for all routes advertised from a unique customer edge (CE) peer or route.
per-vrf	Specifies that the same label is used for all routes advertised from a unique VRF.
per-prefix	Specifies that the same label is used for all routes advertised from a unique prefix. **Note** This keyword is applicable only for IPv4/IPv6 AF configuration mode under vrf mode.
route-policy	Specifies a route policy to select prefixes for setting the label mode.

Command Default

Per-prefix label mode.

Note If a policy attached at label-mode attachpoint evaluates to pass and a **label mode** is not explicitly set, **per-prefix** is used as the default label mode.

If a policy attached at label-mode attachpoint evaluates to a drop, **per-prefix** is used as a default label mode. If any **label mode** is set explicitly in this case, it will be ignored.

Command Modes

VPNv4 unicast address family configuration

VPNv4 address family configuration

VPNv6 address family configuration

VRF IPv4 address family configuration

VRF IPv6 address family configuration

Command History

Release	Modification
Release 6.0	This command was introduced.
Release 7.4.1	This command was deprecated in the VRF IPv4, VRF IPv6, VPNv4 unicast, and VPNv6 unicast address-family configuration modes. The function of this deprecated command can be carried out using **label mode** command under address-family mode.

Usage Guidelines

To configure label mode at VPN-AF level and to have all the VRF AFs inherit that configuration, you must use **vrf all**, which is available under VPN-AF mode.

The inheritance rules followed are:

- **label mode** configuration under VRF-AF, overrides **label-allocation-mode** configuration under VRF and **label mode** configuration under VPN-AF.
- **label-allocation-mode** configuration under VRF, overrides **label mode** configuration under VPN-AF.
- The order of priority to determine the label mode in the configurations is:
  1. VRF-AF: **label mode**
  2. VRF: **label-allocation-mode**
  3. VPN-AF: **label mode**
  4. N/A: **per-prefix**



Note Even if **label mode** is in use, **per-vrf** label is allocated for connected, aggregate, and local prefixes.

The per-nexthop-received-label mode falls back to per-prefix label mode in the following scenarios:

- When there is a best path and a backup path, and one of these paths is received from an external BGP (eBGP) peer.
- When the best-external path feature is configured using the advertise best-external command.
- When an Autonomous System Boundary Router (ASBR) receives remote VPN routes with a Route Distinguisher (RD) (for example, RD1), and the ASBR is also functioning as a local PE router with a local VRF configured for the same RD1.

When configuring the **label mode per-nexthop-received-label** command at the VPN level, you must explicitly configure the local retention time for the **retain local-label** command.

Task ID

Task ID	Operation
bgp	read, write

The example shows how to configure label mode selection at VPNv4 and VPNv6 AF level:

```
Router# configure
Router(config)# router bgp 100
Router(config-bgp)# address-family vpnv4 unicast
Router(config-bgp-af)# label mode per-nexthop-received-label allocate-secondary-label
```

The example shows how to configure label mode selection at VPNv4 AF level under **vrf all** mode:

```
Router# configure
Router(config)# router bgp 100
Router(config-bgp)# address-family vpnv4 unicast
Router(config-bgp-af)# vrf all
Router(config-bgp-af)# label mode route-policy policy_A
```

The example shows how to configure label mode selection at VRF IPv4 AF level:

```
Router# configure
Router(config)# router bgp 109
Router(config-bgp)# vrf vrf-1
Router(config-bgp-vrf)# rd 1:1
Router(config-bgp-vrf)# address-family ipv4 unicast
Router(config-bgp-vrf)# label mode route-policy policy_B
```

local-as

To allow customization of the autonomous system number for external Border Gateway Protocol (eBGP) neighbor peerings, use the **local-as** command in an appropriate configuration mode. To disable customization of local autonomous system values for eBGP neighbor peerings, use the **no** form of this command.

local-as {*as-number* [**no-prepend** [**replace-as** [**dual-as**]]] | **inheritance-disable**}
no local-as [*as-number* [**no-prepend** [**replace-as** [**dual-as**]]] | **inheritance-disable**]

Syntax Description

as-number	Valid autonomous system number. Range for 2-byte Autonomous system numbers (ASNs) is 1 to 65535. Range for 4-byte Autonomous system numbers (ASNs) in asplain format is 1 to 4294967295. Range for 4-byte Autonomous system numbers (ASNs) is asdot format is 1.0 to 65535.65535. Cannot be the autonomous system number to which the neighbor belongs.
no-prepend	(Optional) Specifies that local autonomous system values are not prepended to announcements from the neighbor.
replace-as	(Optional) Specifies that prepend only local autonomous system values to announcements to the neighbor.
dual-as	(Optional) Dual-AS mode.
inheritance-disable	Prevents local AS from being inherited from the parent.

Command Default

The BGP autonomous system number specified in the **router bgp** command is used, except when confederations are in use. The confederation autonomous system is used for external neighbors in an autonomous system that is not part of the confederation.

Command Modes

Neighbor configuration

VRF neighbor configuration

Neighbor group configuration

Session group configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

You can specify the autonomous system number the local BGP uses to peer with each neighbor. The autonomous system number specified with this command cannot be the local BGP autonomous system number (specified with the **router bgp** command) or the autonomous system number of the neighbor (specified with the **remote-as** command). However, from Release 5.2.2, the autonomous system number for **local-as** and **remote-as** can be

the same, which makes the resulting neighbor peering being treated as iBGP. This command cannot be specified for internal neighbors or for external neighbors in an autonomous system that is part of a confederation.

If this command is configured for a neighbor group or session group, all neighbors using the group inherit the configuration. Values of commands configured specifically for a neighbor override inherited values.

In case the router is an ASBR and adding the **local-as** configuration changes the BGP neighbor from being an eBGP to iBGP, incremental addition of the **local-as** command is not supported. Remove the neighbor configuration, then configure the entire neighbor, including **remote-as** configuration and the new **local-as** configuration, in one commit.

Task ID

Task ID	Operations
bgp	read, write

Examples

The following example shows BGP using autonomous system 30 for the purpose of peering with neighbor 172.20.1.1:

```
RP/0/RP0/CPU0:router(config)# router bgp 140
RP/0/RP0/CPU0:router(config-bgp)# neighbor 172.20.1.1
RP/0/RP0/CPU0:router(config-bgp-nbr)# remote-as 300
RP/0/RP0/CPU0:router(config-bgp-nbr)# local-as 30
```

long-lived-graceful-restart

To enable long lived graceful restart (LLGR) on the BGP neighbors, use the **long-lived-graceful-restart** command in neighbor VPN address family mode. To disable LLGR, use the **no** form of this command.

long-lived-graceful-restart { **capable** | **stale-time send** *time* **accept** *time* }

Syntax Description

capable	Treats the neighbor as LLGR capable even if it does not advertise the capabilities.
stale-time	Causes the local router to advertise the LLGR capability to the neighbor and to enable LLGR for prefixes received from the neighbor.
send *time*	Specifies stale-time sent in LLGR capability.
accept*time*	Specifies maximum stale-time acceptable from neighbor.

Command Default

The default send and accept time is zero.

Command Modes

VPNv4 address family configuration

VPNv6 address family configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

When this command is configured, the BGP session is reset, because the changes need to be advertised to the neighbor in a BGP OPEN message.

When the BGP session to a neighbor goes down the routes received from it will be marked LLGR stale if all of the following conditions are met:

- Either the neighbor is configured as capable or the neighbor sent the LLGR capability in its BGP OPEN message.
- The accept time is not configured to be 0.
- The stale time that the neighbor sent in the LLGR capability in its BGP OPEN message is not 0.
- The neighbor session was not brought down with a clear command on the local router.
- The neighbor sent either the LLGR or graceful restart capability in its BGP OPEN message.

LLGR routes will only be advertised to a neighbor that is LLGR capable, either because it is configured as capable or because it has sent the LLGR capability in its BGP OPEN message. An LLGR route is either one that has been marked as LLGR stale, because the BGP session from which it was received went down or because it has the *LLGR_STALE* community and does not have the *NO_LLGR* community.

Task ID

Task ID	Operation
bgp	read

Examples

This example shows how to configure the neighbor to be LLGR capable for the given address family:

```
Router# configure
Router(config)# router bgp 100
Router(config-bgp)# neighbor 3.3.3.3
Router(config-bgp-nbr)# address-family vpnv4 unicast
Router(config-bgp-nbr-af)# long-lived-graceful-restart capable
```

The **long-lived-graceful-restart** capable command enables the LLGR capability on the neighbor; even though the neighbor does not advertise the LLGR capabilities during session information.

The following example shows how to advertise :

```
Router# configure
Router(config)# router bgp 100
Router(config-bgp)# neighbor 3.3.3.3
Router(config-bgp-nbr)# address-family vpnv4 unicast
Router(config-bgp-nbr-af)# long-lived-graceful-restart stale-time send 20 accept 30
```

The **long-lived-graceful-restart stale-time send 20 accept 30** command is used to configure the LLGR on the neighbor. When this command is configured the configured device will retain routes from the neighbor.

match flow-tag

To identify specific flow-tag values as match criteria in a class-map, use the **match flow-tag** command in class-map configuration mode. To remove a specific flow tag value from the matching criteria for a class-map, use the **no** form of this command.

match flow-tag {*flow-tag number* | [*min-value* – *max-value*]}
no match flow-tag {*flow-tag number* | [*min-value* – *max-value*]}

Syntax Description

flow-tag number	A flow-tag number. Range is from 1 to 63.
min-value	Lower limit of the flow-tag to match. Value range is 1 to 63.
max-value	Upper limit of the flow-tag to match. Value range is 1 to 63.

Command Default

No match criteria is specified.

Command Modes

Class-map configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

Task ID

Task ID	Operation
class-map	read, write

Flow-tag for a route

This example uses the**show route** command to display the flow-tag for a given route.

```
RP/0/0/CPU0:ios-xr#
RP/0/0/CPU0:ios-xr#show route 4.4.4.0/24 detail
Thu Nov 14 15:32:21.010 PST
Routing entry for 4.4.4.0/24
Known via "bgp 200", distance 20, metric 0
Tag 300, type external
Installed Nov 14 09:36:55.066 for 05:55:26
Routing Descriptor Blocks
3.3.3.3, from 3.3.3.3, BGP external
Route metric is 0
Label: None
Tunnel ID: None
Extended communities count: 0
NHID:0x0(Ref:0)
Route version is 0x1 (1)
No local label
IP Precedence: Not Set
QoS Group ID: Not Set
```

```
Flow-tag: 220
Route Priority: RIB_PRIORITY_RECURSIVE (12) SVD Type
RIB_SVD_TYPE_LOCAL
Download Priority 4, Download Version 7
No advertising protos.
RP/0/0/CPU0:ios-xr#
```

maximum-paths (BGP)

To control the maximum number of parallel routes that Border Gateway Protocol (BGP) installs in the routing table, use the **maximum-paths** command in an appropriate configuration mode. To set the maximum number of parallel routes the software installs to the default value, use the **no** form of this command.

maximum-paths { { **ebgp** | **ibgp** | **eibgp** } *maximum* [**unequal-cost**] }

no maximum-paths { { **ebgp** | **ibgp** | **eibgp** } *maximum* [**unequal-cost**] }

Syntax Description

ebgp	Specifies external BGP multipath peers.
ibgp	Specifies internal BGP multipath peers.
eibgp	Specifies internal and external BGP multipath peers. eiBGP allows simultaneous use of internal and external paths.
maximum	Maximum number of parallel routes that BGP installs in the routing table. For VRF, this range is 2 to 128. For non-VRF, this range is 2 to 1024.
unequal-cost	(Optional) Allows iBGP multipaths to have different BGP next-hop Interior Gateway Protocol (IGP) metrics. This option is available when the **ibgp** keyword is used.

Command Default

One path is installed in the routing table.

Command Modes

IPv4 address family configuration

IPv6 address family configuration

VRF IPv4 address family configuration

VRF IPv6 address family configuration

Command History

Release	Modification
Release 7.11.1	The range of the maximum number of parallel routes that BGP installs in the routing table for VRF is changed to 2-128.
Release 6.0	This command was introduced.

Usage Guidelines

Use the **maximum-paths** command to allow the BGP protocol to install multiple paths into the routing table for each prefix. Multiple paths are installed for external peers that are from the same autonomous system and are equal cost (according to the BGP best-path algorithm). Similarly, multiple paths are installed for internal peers that are equal cost based on the BGP best-path algorithm. The IGP metric to the BGP next hop is the same as the best-path IGP metric unless the router is configured for unequal cost iBGP multipath or eiBGP multipath. See *Implementing BGP* in the *BGP Configuration Guide for Cisco NCS 5500 Series Routers* for information on the BGP best-path algorithm.



Note The **maximum-paths** command with the **eibgp** keyword cannot be configured if the **ibgp** or **ebgp** keywords have been configured, because the **eibgp** keyword is a superset of the **ibgp** or **ebgp** keywords.

Task ID

Task ID	Operations
bgp	read, write

Examples

The following example shows how to allow a maximum of four paths to a destination installed into the IPv4 unicast routing table:

```
Router(config)# router bgp 109
Router(config)#(config-bgp)# address-family ipv4 unicast
Router(config)#(config-bgp-af)# maximum-paths ebgp 4
```

maximum-paths external-rtfilter (BGP IPv4 RT-filter address family)

To enable eBGP RT-constraint multipath and allow the router to select up to the specified number of external RT-filter paths, use the **maximum-paths external-rtfilter** command in the IPv4 RT-constraint address family configuration mode. To disable this configuration, use the **no** form of this command.

maximum-paths external-rtfilter *path-count*

Syntax Description

path-count	Specifies the number of external RT-filter paths. The path count ranges from 2 to 1024.

Command Default

None

Command Modes

IPv4 RT-constraint address family configuration mode

Command History

Release	Modification
Release 26.3.1	This command was introduced.

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Task ID

Task ID	Operations
bgp	read, write

Examples

This example enables eBGP RT-constraint multipath and allows the router to select up to 240 qualified external RT-filter paths.

```
Router(config)# router bgp 140
Router(config-bgp)# address-family ipv4 rt-filter
Router(config-bgp-af)# maximum-paths external-rtfilter 240
```

To disable the configuration, use the **no maximum-paths external-rtfilter** command.

maximum-prefix (BGP)

To control how many prefixes can be received from a neighbor, use the **maximum-prefix** command in an appropriate configuration mode. To set the prefix limits to the default values, use the **no** form of this command.

maximum-prefix *maximum* [*threshold*] [**discard-extra-paths**] [**warning-only**] [**restart** *time-interval*]
no maximum-prefix *maximum* [*threshold*] [**discard-extra-paths**] [**warning-only**] [**restart** *time-interval*]

Syntax Description

maximum	Maximum number of prefixes allowed from this neighbor. Range is from1 to 4294967295. **Note** When using additional-paths feature, each path with a unique path ID received from a peer is counted separately for the purpose of maximum-prefix functionality. Hence, the *maximum* value should be configured appropriately when the peer is capable of sending additional-paths.
discard-extra-paths	(Optional) Drops all the excess prefixes received from the neighbor when the prefixes exceed the configured maximum value.
threshold	(Optional) Integer specifying at what percentage of the *maximum* argument value the software starts to generate a warning message. Range is from1 to 100.
warning-only	(Optional) Instructs the software to only generate a log message when the *maximum* argument value is exceeded, and not to terminate the peering.
restart *time-interval*	(Optional) Sets the time interval (in minutes) after which peering session should be reestablished. Configure restart time interval in minutes. Range is from 1 to 65535.

Command Default

The default threshold, when a warning message is generated, is 75 percent.

Command Modes

IPv4 address family group, neighbor address family, and neighbor group address family configuration

IPv6 address family group, neighbor address family, and neighbor group address family configuration

IPv4 tunnel address family group, neighbor group address family, and neighbor address family configuration

IPv4 flowspec under neighbor address family, neighbor group address family, and address family group configuration

IPv6 flowspec under neighbor address family, neighbor group address family, and address family group configuration

VPNv4 flowspec under neighbor address family, neighbor group address family, and address family group configuration

VPNv6 flowspec under neighbor address family, neighbor group address family, and address family group configuration

L2VPN EVPN under neighbor address family, neighbor group address family, and address family group configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

Use the **maximum-prefix** command to configure a maximum number of prefixes that a BGP router is allowed to receive from a neighbor. It adds another mechanism (besides routing policy) to control prefixes received from a peer.

When the number of received prefixes exceeds the maximum number configured, the software terminates the peering, by default, after sending a cease notification to the neighbor. However, if the **warning-only** keyword is configured, the software writes only a log message, but continues peering with the sender. If the peer is terminated, the peer stays down until the **clear bgp** command is issued or the **restart** *time-interval* option is used.

This command takes effect immediately if configured on an established neighbor, unless the number of prefixes received from the neighbor already exceeds the configured limits.

If this command is configured for a neighbor group or neighbor address family group, all neighbors using the group inherit the configuration. Values of commands configured specifically for a neighbor override inherited values.

Task ID

Task ID	Operations
bgp	read, write

Examples

This example shows the maximum number of IP Version 6 (IPv6) unicast prefixes allowed from neighbor 192.168.40.25 set to 5000, threshold value 80%, and restart time interval 20 minutes:

```
RP/0/RP0/CPU0:router(config)#router bgp 100
RP/0/RP0/CPU0:router(config-bgp)#neighbor 192.168.40.25
RP/0/RP0/CPU0:router(config-bgp-nbr)#remote-as 1
RP/0/RP0/CPU0:router(config-bgp-nbr)#address-family ipv6 unicast
```

```
RP/0/RP0/CPU0:router(config-bgp-nbr-af)#maximum-prefix 5000 80 restart 20
```

This example shows the maximum number of IP Version 4 (IPv4) unicast prefixes allowed from the neighbor 192.168.40.24 set to 1000:

```
RP/0/RP0/CPU0:router(config-bgp)# router bgp 109
RP/0/RP0/CPU0:router(config-bgp)# neighbor 192.168.40.24
RP/0/RP0/CPU0:router(config-bgp-nbr)# remote-as 1
RP/0/RP0/CPU0:router(config-bgp-nbr)#address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-bgp-nbr-af)# maximum-prefix 1000
```

The following example shows how to configure discard extra paths:

```
RP/0/RP0/CPU0:router#configure
RP/0/RP0/CPU0:router(config)#router bgp 10
RP/0/RP0/CPU0:router(config-bgp)#neighbor 10.0.0.1
RP/0/RP0/CPU0:router(config-bgp-nbr)#address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-bgp-nbr-af)#maximum-prefix 5000 discard-extra-paths
```

maximum-prefix-restart-time (BGP)

To configure the router to automatically reestablish a BGP neighbor session that has been disabled because the maximum-prefix limit has been exceeded, use the **maximum-prefix-restart-time** command in XR Config mode. To set the prefix limits to the default values, use the **no** form of this command.

maximum-prefix-restart-time *seconds*

Syntax Description

seconds	The configurable range of seconds is from 1 to 65535 seconds..

Command Modes

XR Config mode

Command History

Release	Modification
Release 7.10.1	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

Task ID

Task ID	Operation
bgp	read, write

This example shows how to configure **maximum-prefix-restart-time** interval for neighbor 192.168.40.24:

```
Router(config)#router bgp 109
Router(config-bgp)#neighbor 192.168.40.24
Router(config-bgp-nbr)#maximum-prefix-restart-time 4545
Router(config-bgp-nbr)#commit
```

mpls activate (BGP)

To enable Multiprotocol Label Switching (MPLS) on an interface basis for ASBR and CSC configurations whenever a bgp confederation configuration is used, use the **mpls activate** command in bgp configuration mode. This is needed for InterAS (option B and C) and Carrier Supporting Carrier (CSC) configurations with confederations.

The normal InterAS and CSC configurations (without confederations) do not need to enable this.

To restore the system to its default condition, use the **no** form of this command.

mpls activate *interface id*
no mpls activate *interface id*

Syntax Description

interface id	Name of the interface.

Command Default

No default behavior or values

Command Modes

Router configuration

Neighbor configuration

IPv4 address family group configuration

VPNv4 address family group configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

The**mpls activate** command enables MPLS on the interface specified and also adds the implicit null rewrite corresponding to the peer associated with the interface. The interface specified must be the one corresponding to the inter-AS ASBR or CSC peer.

Task ID

Task ID	Operations
bgp	read, write

Examples

The following example shows how to activate MPLS for InterAS Option B (with confederations):

```
RP/0/RP0/CPU0:router(config)#router bgp 1

  bgp confederation peers

     2002

    !
```

```
 bgp confederation identifier 4589
 bgp router-id 3.3.3.3
 mpls activate
  interface GigabitEthernet0/1/0/0
 !
 address-family ipv4 unicast
  redistribute connected
 !
 address-family vpnv4 unicast
  retain route-target all
 !
 neighbor 10.0.0.9
  remote-as 2002
  address-family ipv4 unicast
   route-policy pass in
   route-policy pass out
  !
  address-family vpnv4 unicast
   route-policy pass in
```

The following example shows how to activate MPLS for CSC (with confederations):

```
router bgp 2002
  bgp confederation peers
   1
  !
  bgp confederation identifier 4589
  bgp router-id 4.4.4.4
  address-family ipv4 unicast
   allocate-label all
  !
  address-family vpnv4 unicast
   retain route-target all
  !
  vrf foo
   rd 1:1
   mpls activate
    interface GigabitEthernet0/1/0/2
   !
   address-family ipv4 unicast
    redistribute connected
    allocate-label all
   !
```

```
    neighbor 10.0.0.1
     remote-as 1
     address-family ipv4 unicast
     !
     address-family ipv4 labeled-unicast
      route-policy pass in
      route-policy pass out
     !
    !
   !
  !
RP/0/RP0/CPU0:router#show mpls forwarding
  Local  Outgoing    Prefix             Outgoing Next Hop        Bytes
  Label  Label       or ID              Interface
  Switched
  ------ ----------- ------------------ ------------ ---------------
  ------------
  16000  Aggregate   foo: Per-VRF Aggr[V]   \
                                           foo                        0
  16001  Pop         10.0.0.0/16[V]     Gi0/1/0/2    10.0.0.1        44



RP/0/RP0/CPU0:router#show mpls interfaces
  Interface                  LDP      Tunnel   Enabled
  -------------------------- -------- -------- --------
  GigabitEthernet0/1/0/2     No       No       Yes
```

mvpn

To enable BGP instance to connect to PIM/PIM6, use the **mvpn** command in XR Config mode. To disable BGP instance -PIM/PIM6 connection, use the **no** form of this command.

mvpn
no mvpn

Syntax Description

This command has no keywords or arguments.

Command Default

PIM/PIM connection is disabled.

Command Modes

XR Config mode

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

Task ID

Task ID	Operation
bgp	read, write

This example shows how to configure mvpn and enable PIM/PIM6 connection:

```
RP/0/RP0/CPU0:router#configure
RP/0/RP0/CPU0:router(config)#router bgp 100
RP/0/RP0/CPU0:router(config-bgp)#mvpn
```

neighbor (BGP)

To enter neighbor configuration mode for configuring Border Gateway Protocol (BGP) routing sessions, use the **neighbor** command in an appropriate configuration mode. To delete all configuration for a neighbor and terminate peering sessions with the neighbor, use the **no** form of this command.

neighbor *ip-address*
no neighbor *ip-address*

Syntax Description

ip-address	IPv4 or IPv6 IP address of the BGP-speaking neighbor.

Command Default

Neighbor mode is not specified.

Command Modes

Router configuration

VRF configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

From router configuration mode, you can use this command to enter neighbor configuration mode.

From neighbor configuration mode, you can enter address family configuration for the neighbor by using the **address-family** command, which allows you to configure routing sessions for IP Version 4 and IP Version 6 address prefixes.

The **neighbor** command does not cause the neighbor to be configured and does not result in a peering to be established with the neighbor. To create the neighbor, you configure a remote autonomous system number by entering the **remote-as** command, or the neighbor can inherit a remote autonomous system from a neighbor group or session group if the **use** command is applied.



Note A neighbor must have must a remote autonomous system number, and an IP address and address family must be enabled on the neighbor.

Unlike IPv4, IPv6 must be enabled before any IPv6 neighbors can be defined. Enable IPv6 in router configuration mode using the **address-family** command.



Note Configuration for the neighbor cannot occur (peering is not established) until the neighbor is given a remote as-number and neighbor address.

The **no** form of this command causes the peering with the neighbor to be terminated and all configuration that relates to the neighbor to be removed.

Task ID

Task ID	Operations
bgp	read, write

Examples

The following example shows how to place the router in neighbor configuration mode for BGP routing process 1 and configure the neighbor IP address 172.168.40.24 as a BGP peer:

```
RP/0/RP0/CPU0:router(config)# router bgp 1
RP/0/RP0/CPU0:router(config-bgp)# neighbor 172.168.40.24
RP/0/RP0/CPU0:router(config-bgp-nbr)# remote-as 65000
```

The following example shows how to enable IPv6 for BGP, then place the router in neighbor configuration mode for an IPv6 neighbor, 3000::1, and configure neighbor 3000::1 as a BGP peer:

```
RP/0/RP0/CPU0:router(config)# router bgp 100
RP/0/RP0/CPU0:router(config-bgp)# address-family ipv6 unicast
RP/0/RP0/CPU0:router(config-bgp-af)# exit
RP/0/RP0/CPU0:router(config-bgp)# neighbor 3000::1
RP/0/RP0/CPU0:router(config-bgp-nbr)# remote-as 2002
RP/0/RP0/CPU0:router(config-bgp-nbr)# address-family ipv6 unicast
```

neighbor-group

To create a neighbor group and enter neighbor group configuration mode, use the **neighbor-group** command in XR Config mode. To remove a neighbor group and delete all configuration associated with the group, use the **no** form of this command.

neighbor-group *name*
no neighbor-group *name*

Syntax Description

name	Neighbor group name.

Command Default

No neighbor group mode is specified.

Command Modes

XR Config mode

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

The **neighbor-group** command puts the router in neighbor group configuration mode and creates a neighbor group.

A neighbor group helps you apply the same configuration to one or more neighbors. After a neighbor group is configured, each neighbor can inherit the configuration through the **use** command. If a neighbor is configured to use a neighbor group, the neighbor, by default, inherits the entire configuration of the neighbor group, which includes the address family-independent and address family-specific configurations. The inherited configuration can be overridden if you directly configure commands for the neighbor or if you configure session groups or address family groups with the **use** command.

From neighbor group configuration mode, you can configure address family-independent parameters for the neighbor group. To enter address family-specific configuration for the neighbor group, use the **address-family** command when in the neighbor group configuration mode.

Note If an address family is configured for a neighbor group, neighbors that use the neighbor group attempt to exchange routes in that address family.

The **no** form of this command ordinarily causes all configuration for the neighbor group to be removed. If using the **no** form would result in a neighbor losing its remote autonomous system number, the configuration is rejected. In this scenario, the neighbor configuration must be either removed or configured with a remote autonomous system number before the neighbor group configuration can be removed.

Note Neighbor groups should not be configured with a mixture of IPv4 and IPv6 address families, because such a neighbor group is not usable by any neighbor. Note that within the Cisco IOS XR system configuration architecture, it is possible to create such a neighbor group; however, any attempt to use it is rejected.

Task ID

Task ID	Operations
bgp	read, write

Examples

The following example shows how to create a neighbor group called group1 that has IP Version 4 (IPv4) unicast and IPv4 multicast activated along with various configuration features. The neighbor group is used by neighbor 10.0.0.1 and neighbor 10.0.0.2, which allows them to inherit the entire group1 configuration.

```
RP/0/RP0/CPU0:router(config)# router bgp 65530
RP/0/RP0/CPU0:router(config-bgp)# neighbor-group group1
RP/0/RP0/CPU0:router(config-bgp-nbrgrp)# remote-as 65535
RP/0/RP0/CPU0:router(config-bgp-nbrgrp)# advertisement-interval 2
RP/0/RP0/CPU0:router(config-bgp-nbrgrp)# address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-bgp-nbrgrp-af)# send-community-ebgp
RP/0/RP0/CPU0:router(config-bgp-nbrgrp-af)# exit
RP/0/RP0/CPU0:router(config-bgp-nbrgrp)# address-family ipv4 multicast
RP/0/RP0/CPU0:router(config-bgp-nbrgrp-af)# next-hop-self
RP/0/RP0/CPU0:router(config-bgp-nbrgrp-af)# exit
RP/0/RP0/CPU0:router(config-bgp-nbrgrp)# exit
RP/0/RP0/CPU0:router(config-bgp)#neighbor 10.0.0.1
RP/0/RP0/CPU0:router(config-bgp-nbr)# use neighbor-group group1
RP/0/RP0/CPU0:router(config-bgp-nbr)# exit
RP/0/RP0/CPU0:router(config-bgp)# neighbor 10.0.0.2
RP/0/RP0/CPU0:router(config-bgp-nbr)# use neighbor-group group1
RP/0/RP0/CPU0:router(config-bgp-nbr)# exit
```

neighbor internal-vpn-client

To preserve the iBGP-CE (customer edge) attributes inside the VPN attribute set (ATTR-SET) and send it across to the core, use the **neighbor internal-vpn-client** command in the VRF neighbor configuration mode. To disable the command, use the **no** form of this command.

neighbor *ip-address* **internal-vpn-client**

no neighbor *ip-address* **internal-vpn-client**

Syntax Description

neighbor *ip-address*	IP address of the neighboring device.
internal-vpn-client	Stacks the iBGP-CE neighbor path in the VPN attribute set.

Command Default

None

Command Modes

VRF neighbor configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

The **neighbor** ip-address **internal-vpn-client** command enables PE devices to make the entire VPN cloud act as an internal VPN client to the CE devices connected internally. This command is used so that existing internal BGP VRF lite scenarios are not affected. You need not configure autonomous system override for CE devices after enabling this command.

Task ID

Task ID	Operations
bgp	read, write

Examples

The following example shows how to configure L3VPN iBGP PE-CE:

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router bgp 100
RP/0/RP0/CPU0:router(config-bgp)# vrf blue neighbor 10.10.10.1
RP/0/RP0/CPU0:router(config-bgp-vrf-nbr)# internal-vpn-client
```

network (BGP)

To specify that the Border Gateway Protocol (BGP) routing process should originate and advertise a locally known network to its neighbors, use the **network** command in an appropriate configuration mode. To disable originating or advertising the network to neighbors, use the **no** form of this command.

network {*ip-address/prefix-length ip-address mask*} [**route-policy** *route-policy-name*]
no network {*ip-address/prefix-length ip-address mask*} [**route-policy** *route-policy-name*]

Syntax Description

ip- *address*	Network that BGP advertises.
/ *prefix-length*	Length of the IP address prefix. A decimal value that indicates how many of the high-order contiguous bits of the address compose the prefix (the network portion of the address). A slash (/) must precede the decimal value.
ip-address mask	Network mask applied to the *ip-address* argument.
route-policy *route-policy-name*	(Optional) Specifies a route policy to use to modify the attributes of the network.

Command Default

No networks are specified.

Command Modes

IPv4 address family configuration

IPv6 address family configuration

VRF IPv4 address family configuration

VRF IPv6 address family configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

A network specified with this command is originated and advertised to neighbors only if there exists a route for the network in the routing table. That is, there must be a route learned using local or connected networks, static routing, or a dynamic IGP such as IS-IS or OSPF.

Other than the available system resources on the router, no limit exists on the number of network commands that can be configured.

Task ID

Task ID	Operations
bgp	read, write

Examples

The following example shows how to configure the local router to originate the IPv4 unicast network 172.20.0.0/16:

```
RP/0/RP0/CPU0:router(config)#router bgp 120
RP/0/RP0/CPU0:router(config-bgp)# address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-bgp-af)# network 172.20.0.0/16
```

network backdoor

To set the administrative distance on an external Border Gateway Protocol (eBGP) route to that of a locally sourced BGP route, causing it to be less preferred than an Interior Gateway Protocol (IGP) route, use the **network backdoor** command in an appropriate configuration mode. To disable setting the administrative distance to the value for locally sourced BGP routes, use the **no** form of this command.

network {*ip-address/prefix-length ip-address mask*} **backdoor**
no network {*ip-address/prefix-length ip-address mask*} **backdoor**

Syntax Description

ip-address	Network that provides a backdoor route.
/ *prefix-length*	Length of the IP address prefix. A decimal value that indicates how many of the high-order contiguous bits of the address compose the prefix (the network portion of the address). A slash (/) must precede the decimal value.
mask	Network mask applied to the *ip-address* argument.

Command Default

No backdoor routes are installed.

Command Modes

IPv4 address family configuration

IPv6 address family configuration

VRF IPv4 address family configuration

VRF IPv6 address family configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

Configuring the **network backdoor** command does not cause BGP to originate a network, even if an IGP route for the network exists. Ordinarily, the backdoor network would be learned through both an eBGP and IGP. The BGP best-path selection algorithm does not change when a network is configured as a backdoor network.

Task ID

Task ID	Operations
bgp	read, write

Examples

The following example shows IP Version 4 (IPv4) unicast network 192.168.40.0/24 configured as a backdoor network:

```
RP/0/RP0/CPU0:router(config)# router bgp 109
RP/0/RP0/CPU0:router(config-bgp)# address-family ipv4 unicast
```

```
RP/0/RP0/CPU0:router(config-bgp-af)# network 192.168.40.0/24 backdoor
```

next-hop-self

To disable next-hop calculation and insert your own address in the next-hop field of Border Gateway Protocol (BGP) updates, use the **next-hop-self** command in an appropriate configuration mode. To enable next-hop calculation, use the **no** form of this command.

next-hop-self [**inheritance-disable**]
no next-hop-self [**inheritance-disable**]

Syntax Description

inheritance-disable	(Optional) Allows a next-hop calculation override when this feature may be inherited from a neighbor group or address family group.

Command Default

When this command is not specified, the software calculates the next hop for BGP updates accepted by the router.

Command Modes

IPv4 address family group configuration

IPv6 address family group configuration

IPv4 neighbor address family configuration

IPv4 neighbor group address family configuration

IPv6 neighbor group address family configuration

VPNv4 address family group configuration

VPNv4 neighbor address family configuration

VPNv4 neighbor group address family configuration

VPNv6 neighbor group address family configuration

VPNv6 neighbor address family configuration

IPv4 labeled-unicast address family configuration

IPv6 labeled-unicast address family configuration

VRF labeled-unicast address family configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

Use the **next-hop-self** command to set the BGP next-hop attribute of routes being advertised over a peering session to the local source address of the session.

This command is useful in nonmeshed networks in which BGP neighbors may not have direct access to all other neighbors on the same IP subnet.

If this command is configured for a neighbor group or address family group, a neighbor using the group inherits the configuration. Configuring the command specifically for a neighbor overrides any inherited value.

Configuring the **next-hop-self** command under IPv4 labeled-unicast, IPv6 labeled-unicast, or VRF labeled-unicast address family configuration mode enables next-hop-self for labeled prefixes advertised to an iBGP peer.

Task ID

Task ID	Operations
bgp	read, write

Examples

The following example shows how to set the next hop of the update field for all IP Version 4 (IPv4) unicast routes advertised to neighbor 172.20.1.1 to an address of the local router:

```
RP/0/RP0/CPU0:router(config)# router bgp 140
RP/0/RP0/CPU0:router(config-bgp)# neighbor 172.20.1.1
RP/0/RP0/CPU0:router(config-bgp-nbr)# remote-as 1
RP/0/RP0/CPU0:router(config-bgp-nbr)# address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-bgp-nbr-af)# next-hop-self
```

The following example shows how to disable the **next-hop-self** command for neighbor 172.20.1.1. If not overridden, the next hop would be inherited from address family group group1:

```
RP/0/RP0/CPU0:router(config)# router bgp 140
RP/0/RP0/CPU0:router(config-bgp)# af-group group1 address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-bgp-afgrp)# next-hop-self
RP/0/RP0/CPU0:router(config-bgp-afgrp)# exit
RP/0/RP0/CPU0:router(config-bgp)# neighbor 172.20.1.1
RP/0/RP0/CPU0:router(config-bgp-nbr)# remote-as 1
RP/0/RP0/CPU0:router(config-bgp-nbr)# address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-bgp-nbr-af)# use af-group group1
RP/0/RP0/CPU0:router(config-bgp-nbr-af)# next-hop-self inheritance-disable
```

next-hop-unchanged

To disable overwriting of the next hop before advertising to external Border Gateway Protocol (eBGP) peers, use the **next-hop-unchanged** command in an appropriate configuration mode. To enable overwriting of the next hop, use the **no** form of this command.

next-hop-unchanged [**inheritance-disable**]
no next-hop-unchanged [**inheritance-disable** | **multipath**]

Syntax Description

inheritance-disable	(Optional) Allows overwriting of the next hop before advertising to eBGP peers when this feature may be inherited from a neighbor group or address family group.
multipath	(Optional) Disables overwriting of next-hop calculation for multipath prefixes.

Command Default

Overwriting of the next hop is allowed.

Command Modes

IPv4 address family configuration

IPv6 address family configuration

VPNv4 address family group configuration

VPNv4 neighbor address family configuration

VPNv4 neighbor group address family configuration

VPNv6 address family group configuration

VPNv6 neighbor address family configuration

VPNv6 neighbor group address family configuration

IPv4 labeled-unicast address family configuration

IPv6 labeled-unicast address family configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

Use the **next-hop-unchanged** command to propagate the next hop unchanged for multihop eBGP peering sessions. This command should not be configured on a route reflector, and the **next-hop-self** command should not be used to modify the next-hop attribute for a route reflector when this feature is enabled for a route reflector client.



Note Incorrectly setting BGP attributes for a route reflector can cause inconsistent routing, routing loops, or a loss of connectivity. Setting BGP attributes for a route reflector should be attempted only by an experienced network operator.

Task ID

Task ID	Operations
bgp	read, write

Examples

The following example shows how to disable the overwriting of next hops before advertising to eBGP peers:

```
RP/0/RP0/CPU0:router(config)# router bgp 140
RP/0/RP0/CPU0:router(config-bgp)# af-group group1 address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-bgp-afgrp)# next-hop-unchanged disable
RP/0/RP0/CPU0:router(config-bgp-afgrp)# exit
```

The following example shows how to disable the overwriting of next hops for multipath prefixes:

```
RP/0/RP0/CPU0:router(config)# router bgp 100
RP/0/RP0/CPU0:router(config-bgp)# af-group group1 address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-bgp-afgrp)# next-hop-unchanged multipath
RP/0/RP0/CPU0:router(config-bgp-afgrp)# exit
```

nexthop resolution prefix-length minimum

To set minimum prefix-length for nexthop resolution, use the **nexthop resolution prefix-length minimum** command in an appropriate configuration mode. To disable the minimum prefix-length for nexthop resolution, use the **no** form of this command.

nexthop resolution prefix-length minimum *prefix-length-value*
no nexthop resolution prefix-length minimum *prefix-length-value*

Syntax Description

prefix-length-value	Sets the minimum prefix-length. Range is 0 to 32.

Command Default

Nexthop resolution for minimum prefix-length is disabled.

Command Modes

VPNv4 Unicast address family

VRF IPv4 Unicast address family

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

Task ID

Task ID	Operation
bgp	read, write

This example shows how to set the minimum prefix-length for nexthop resolution as *32*:

```
RP/0/RP0/CPU0:router#configure
RP/0/RP0/CPU0:router(config)#router bgp 100
RP/0/RP0/CPU0:router(config-bgp)#address-family vpnv4 unicast
RP/0/RP0/CPU0:router(config-bgp-af)#nexthop resolution prefix-length minimum 32
```

nexthop route-policy

To specify that BGP routes are resolved using only next hops whose routes match specific characteristics, use the **nexthop route-policy** command in the appropriate configuration mode. To remove the **nexthop route-policy** command from the configuration file and restore the system to its default behavior, use the **no** form of this command.

nexthop route-policy *route-policy-name*
no nexthop route-policy *route-policy-name*

Syntax Description

route-policy-name	Route policy to use for filtering based on next hops.

Command Default

No default behavior or values

Command Modes

IPv4 address family configuration

IPv6 address family configuration

VPNv4 address family configuration

VPNv6 address family configuration

VRF address family configuration

Command History

Release	Modification
Release 6.0	This command was introduced.
Release 7.11.1	This command was modified. This command is extended to VRF address-family configuration mode.

Usage Guidelines

Use the **nexthop route-policy** command to configure route policy filtering using next hops.

The BGP next-hop tracking feature allows you to specify that BGP routes are resolved using only next hops whose routes have the following characteristics:

- To avoid the aggregate routes, the prefix length must be greater than a specified value.
- The source protocol must be from a selected list, ensuring that BGP routes are not used to resolve next hops that could lead to oscillation.

This route policy filtering is possible because RIB identifies the source protocol of a route that resolves a next hop as well as the mask length associated with the route.

The next-hop attach point supports matching using the protocol name and mask length. BGP marks all next hops that are rejected by the route policy as invalid, and no best path is calculated for the routes that use the invalid next hop. The invalid next hops continue to stay in the active cache and can be displayed as part of the **show bgp nexthop** command with an invalid status.

Task ID

Task ID	Operations
bgp	read, write

Examples

The following example shows how to specify the route policy nexthop_A as the policy to use for filtering next hops:

```
Router(config)# router bgp 109
Router(config-bgp)# address-family ipv4 unicast
Router(config-bgp-af)# nexthop route-policy nexthop_A
```

The following example shows how to specify the route policy nexthop_A as the policy to use for filtering next hops in the VRF address family configuration mode:

```
Router(config)# router bgp 500
Router(config-bgp)# vrf vrf10
Router(config-bgp-vrf)# address-family ipv4 unicast
Router(config-bgp-vrf-af)# nexthop route-policy nh-route-policy
```

nexthop trigger-delay

To specify the delay for triggering next-hop calculations, use the **nexthop trigger-delay** command in the appropriate configuration mode. To set the trigger delay to the default value, use the **no** form of this command.

nexthop trigger-delay {**critical** *delay* | **non-critical** *delay*}
no nexthop trigger-delay {**critical** *delay* | **non-critical** *delay*}

Syntax Description

critical	Specifies critical next-hop events. For example, when the next hop is unreachable.
delay	Trigger delay, in milliseconds. Range is 0 to 4294967295.
non-critical	Specifies noncritical next-hop events. For example, Interior Gateway Protocol (IGP) metric changes.

Command Default

critical : 3000 msec for all address families, except VPNv4 and VPNv6 address families

critical: 50 msec for VPNv4 and VPNv6 address families

non-critical: 10000 msec for all address families

Command Modes

All address families

Command History

Release	Modification
Release 7.10.1	The default value of **critical** delay for VPN address family was changed from 0 msec to 50 msec.
Release 6.0	This command was introduced.

Usage Guidelines

Use the **nexthop trigger-delay** command to allow for a dynamic way for Interior Gateway Protocol (IGP) to converge. This convergence allows BGP to accumulate all notifications and trigger fewer walks, resulting in fewer interprocess communications (IPCs) to the Routing Information Base (RIB) for route addition, deletion, and modification and fewer updates to peers.

Note A high *delay* value can be configured to effectively turn off next-hop tracking.

The **non-critical** *delay* value must always be set to at least equal or greater than the **critical** *delay value* .

The *delay* should be slightly higher than the time it takes for the IGP to settle into a steady state after some event (IGP convergence time).

Avoid configuring the **nexthop trigger-delay critical** *0* as it is not suitable on:

- Scaled deployments where a long BGP next-hop walk time duration is inevitable.
- Deployments where BGP next-hop changes are frequent.

Disadvantages of nexthop trigger-delay critical *0* configuration

- High CPU utilization as each change notification triggers a BGP next-hop walk for address families configured with **nexthop trigger-delay critical** *0*.
- BGP next-hop change notifications are not batched. This disallows interleaving of next-hop walks in address families with the non-zero delay configuration as these address families wait until the address families with the zero critical delay value complete their next-hop walks.
- Extended wait time before the BGP next-hop walk starts on address families with the non-zero critical delay configuration, leading to potential traffic blackholing.

Starting with Cisco IOS XR Release 7.10.1, the default critical delay configuration in VPNv4 address family was changed from 0 msec to 50 msec. With this change, all address families have a default non-zero critical delay value. To see the critical delay value of each address family, run the **show bgp all all nexthops** command.

After you have upgraded to Cisco IOS XR Release 7.10.1 or later, if you configure the default critical delay value in the IPv4 address family to 0 msec, you will observe a considerable delay in VPNv4 convergence for the following reasons:

- The IPv4 address families are walked as many times as the number of next-hop critical alerts raised to BGP.
- The BGP next-hop updates for the IPv4 address family prefixes take precedence over VPNv4 address family prefixes.

Advantages of configuring nexthop trigger-delay critical with a non-zero default value

- Provides next-hop change notification batching which reduces the number of BGP next-hop walks.
- Allows interleaving different active BGP next-hop walks for the respective address families while prioritizing some address families over the others.

Therefore, we strongly recommend you to configure **nexthop trigger-delay critical** with a non-zero value.

Task ID

Task ID	Operations
bgp	read, write

Examples

The following example shows how to set the critical next-hop trigger delay to 3500 milliseconds:

```
RP/0/RP0/CPU0:router(config)# router bgp 109
RP/0/RP0/CPU0:router(config-bgp)# address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-bgp-af)# nexthop trigger-delay critical 3500
```

orf

To specify Outbound Route Filter (ORF) and inbound filtering criteria, use the **orf route-policy** command in an appropriate configuration mode. To restore the system to its default condition, use the **no** form of this command.

orf route-policy *route-policy-name*
no orf route-policy *route-policy-name*

Syntax Description

route-policy-name	Name of the route policy.

Command Default

No ORF route policy is defined.

Command Modes

IPv4 address family group configuration

IPv6 address family group configuration

IPv4 neighbor address family configuration

IPv4 neighbor group address family configuration

IPv6 neighbor group address family configuration

VRF IPv4 neighbor address family configuration

VRF IPv6 neighbor address family configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

Task ID

Task ID	Operations
bgp	read, write

Examples

The following example shows how to configure outbound and inbound filtering criteria:

```
RP/0/RP0/CPU0:router(config)#router bgp 6
RP/0/RP0/CPU0:router(config-bgp)# neighbor 172.20.1.1
RP/0/RP0/CPU0:router(config-bgp-nbr)# remote-as 1
RP/0/RP0/CPU0:router(config-bgp-nbr)# address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-bgp-nbr-af)#orf route-policy policy_A
```

password (BGP)

To enable Message Digest 5 (MD5) authentication on a TCP connection between two Border Gateway Protocol (BGP) neighbors, use the **password** command in an appropriate configuration mode. To disable MD5 authentication, use the **no** form of this command.

password {**clear** | **encrypted**} *password*
no password [**clear** *password* | **encrypted** *password*]

Syntax Description

clear	Specifies that an unencrypted password follows. The password must be a case-sensitive, clear-text unencrypted password.
encrypted	Specifies that an encrypted password follows. The password must be a case-sensitive, encrypted password.
password	Password of up to 80 characters. The password can contain any alphanumeric characters. However, if the first character is a number or the password contains a space, the password must be enclosed in double quotation marks; for example, “2 password.”

Command Default

When this command is not specified in the appropriate configuration mode, MD5 authentication is not enabled on a TCP connection between two BGP neighbors.

Command Modes

Neighbor configuration

VRF neighbor configuration

Neighbor group configuration

Session group configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

Configure a password to enable authentication between two BGP peers. Use the **password** command to verify each segment sent on the TCP connection between the peers. The same password must be configured on both networking devices, otherwise a connection cannot be made. The authentication feature uses the MD5 algorithm. Specifying this command causes the software to generate and check the MD5 digest on every segment sent on the TCP connection.

Configuring a neighbor password does not cause the existing session for a neighbor to end. However, until the new password is configured on the remote router, the local BGP process does not receive keepalive messages from the remote device. If the password is not updated on the remote device by the end of the hold time, the session ends. The hold time can be changed using the **timers** command or the **timers bgp** command.

If this command is configured for a neighbor group or neighbor address family group, a neighbor using the group inherits the configuration. Values of commands configured specifically for a neighbor overrides inherited values.

Task ID

Task ID	Operations
bgp	read, write

Examples

The following example shows how to configure neighbor 172.20.1.1 to use MD5 authentication with the password password1:

```
RP/0/RP0/CPU0:router(config)# router bgp 140
RP/0/RP0/CPU0:router(config-bgp)#neighbor 172.20.1.1
RP/0/RP0/CPU0:router(config-bgp-nbr)#remote-as 1
RP/0/RP0/CPU0:router(config-bgp-nbr)#password clear password1
```

password (rpki-server)

To specify a SSH password for the RPKI cache-server, use the **password** command in rpki-server configuration mode. To remove the SSH passwords, use the **no** form of this command.

password *password*
no password *password*

Syntax Description

password	Enters a password to be used for the SSH transport mechanism.

Command Default

Password is not configured.

Command Modes

RPKI server configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

SSH expects to use an authentication method to connect to a remote server. The SSH authentication method to connect to RPKI server is password-based. So, the RPKI cache-server must be configured with username and password. A username and password must be configure for each server configured under BGP that uses the SSH transport

Task ID

Task ID	Operation
bgp	read, write

This example shows how to configure a username (*rpki-user*) and password (*rpki-ssh-pass*) for the RPKI cache-server SSH transport mechanism:

```
RP/0/RP0/CPU0:router#configure
RP/0/RP0/CPU0:router(config)#router bgp 100
RP/0/RP0/CPU0:router(config-bgp)#rpki server 172.168.35.40
RP/0/RP0/CPU0:router(config-bgp-rpki-cache)# transport ssh port 22
RP/0/RP0/CPU0:router(config-bgp-rpki-cache)#username rpki-user
RP/0/RP0/CPU0:router(config-bgp-rpki-cache)#password rpki-ssh-pass
```

password-disable

To override any inherited password configuration from a neighbor group or session group for Border Gateway Protocol (BGP) neighbors, use the **password-disable** command in an appropriate configuration mode. To disable overriding any inherited password command, use the **no** form of this command.

password-disable
no password-disable

Syntax Description

This command has no arguments or keywords.

Command Default

Configured passwords for neighbor and session groups are inherited.

Command Modes

Neighbor configuration

VRF neighbor configuration

Neighbor group configuration

Session group configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

If you specify a password on a neighbor group or session group, all users of the group inherit the password. Specifying a different **password** command specifically on a neighbor that uses the group overrides the inherited value. Specifying **password-disable** on a neighbor that uses the group disables password authentication for the neighbor.

Task ID

Task ID	Operations
bgp	read, write

Examples

The following example shows how to disable MD5 authentication for neighbor 172.20.1.1, preventing it from inheriting the password password1 from session group group1:

```
RP/0/RP0/CPU0:router(config)# router bgp 140
RP/0/RP0/CPU0:router(config-bgp)# session-group group1
RP/0/RP0/CPU0:router(config-bgp-sngrp)# password clear password1
RP/0/RP0/CPU0:router(config-bgp-sngrp)# exit
RP/0/RP0/CPU0:router(config-bgp)# neighbor 172.20.1.1
RP/0/RP0/CPU0:router(config-bgp-nbr)# remote-as 2
RP/0/RP0/CPU0:router(config-bgp-nbr)# use session-group group1
RP/0/RP0/CPU0:router(config-bgp-nbr)# password-disable
```

permanent-network

To define a prefix set as permanent, use the **permanent-network** command in the global address family configuration mode. To remove a prefix set as permanent, use the **no** form of this command. The **permanent-network** command uses a route-policy to identify the set of prefixes (networks) for which permanent paths needs to be created.

The permanent network feature supports only prefixes in IPv4 unicast and IPv6 unicast address-families under the default Virtual Routing and Forwarding (VRF).

permanent-network route-policy *route-policy-name*
no permanent-network

Syntax Description

route-policy *route-policy-name*	Specifies a configured routing policy.

Command Default

None

Command Modes

Address-family configuration.

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

Task ID

Task ID	Operations
bgp	read, write

Examples

This example shows how to define permanent path for a route policy named POLICY-PERMANENT-NETWORK-IPv4:

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router bgp 100
RP/0/RP0/CPU0:router(config-bgp)# address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-af)# permanent-network route-policy POLICY-PERMANENT-NETWORK-IPv4
```

precedence

To set the precedence level, use the **precedence** command in the appropriate configuration mode. To remove the **precedence** command from the configuration file and restore the system to its default interval values, use the **no** form of this command.

precedence *value*
no precedence [*value*]

Syntax Description

value Value of the precedence. The precedence value can be a number from 0 to 7, or it can be one of the following keywords:

routine —Set packets with routine precedence (0)

priority —Set packets with priority precedence (1)

immediate —Set packets with immediate precedence (2)

flash — Set packets with flash precedence (3)

flash-override —Set packets with flash override precedence (4)

critical —Set packets with critical precedence (5)

internet —Set packets with internetwork control precedence (6)

network —Set packets with network control precedence (7)

Command Default

No default behavior or values

Command Modes

Neighbor configuration

Neighbor session group configuration

Neighbor group configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

Use the **precedence** command to set the precedence value.

Task ID

Task ID	Operations
bgp	read, write

Examples

The following example shows how to set the precedence to 2:

```
RP/0/RP0/CPU0:router(config)# router bgp 5
RP/0/RP0/CPU0:router(config-bgp)# neighbor 10.1.1.1
```

```
RP/0/RP0/CPU0:router(config-bgp-nbr)# remote-as 100
RP/0/RP0/CPU0:router(config-bgp-nbr)# precedence 2
```

preference (rpki-server)

To specify a preference value for the RPKI cache-server, use the **preference** command rpki-server configuration mode. To remove the preference value, use the **no** form of this command.

preference *preference-value*
no preference *preference-value*

Syntax Description

preference-value	Specifies a RPKI cache preference value. Range is 1 to 10. **Note** A lower value is recommended.

Command Default

Preference value is not set.

Command Modes

RPKI server configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

Task ID

Task ID	Operation
bgp	read, write

This example shows how to set preference value for RPKI configuration as 1:

```
RP/0/RP0/CPU0:router#configure
RP/0/RP0/CPU0:router(config)#router bgp 100
RP/0/RP0/CPU0:router(config-bgp)#rpki server 172.168.35.40
RP/0/RP0/CPU0:router(config-bgp-rpki-cache)# transport ssh port 22
RP/0/RP0/CPU0:router(config-bgp-rpki-cache)#username rpki-user
RP/0/RP0/CPU0:router(config-bgp-rpki-cache)#password rpki-ssh-pass
RP/0/RP0/CPU0:router(config-bgp-rpki-cache)#preference 1
```

prefix-ecmp-delay

To configure ECMP delay duration and the resource usage threshold limit in BGP, use the **prefix-ecmp-delay** *interval_value* **oor-threshold** *threshold_value* command in address-family configuration mode.

prefix-ecmp-delay *interval_value* **oor-threshold** *threshold_value*

Syntax Description

interval_value	Specifies the duration at which the route updates are delayed.
threshold_value	Specifies the resource usage threshold.

Command Default

No default behavior or values

Command Modes

IPv4/IPv6 address-family configuration

Command History

Release	Modification
Release 24.2.1	This command was introduced.

Usage Guidelines

None.

Examples

The following example shows how to configure a delay of 10 seconds when the resource usage exceeds a threshold of 30%:

```
router bgp 100
  address-family ipv4 unicast
    prefix-ecmp-delay 10000 oor-threshold 30
```

purge-time (rpki-server)

To configure the time BGP waits to keep routes from RPKI cache-server after the cache session drops, use the **purge-time** command in rpki-server configuration mode. To remove the purge-time configuration, use the **no** form of this command.

purge-time *time-in-seconds*
no purge-time *time-in-seconds*

Syntax Description

time-in-seconds	Sets the purge time in seconds. Range is 30 to 360 seconds.

Command Default

Sets a default value of 60 seconds.

Command Modes

RPKI server configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

When a cache session is dropped then a "purge-timer" is started for that cache. If the session re-establishes within the timer interval, then the purge timer is stopped and no further action is taken. If the cache session does not re-establish within the timer interval, only then does BGP remove all ROAs from the cache.

Task ID

Task ID	Operation
bgp	read, write

This example shows how to set the purge-time for RPKI cache as 30 seconds:

```
RP/0/RP0/CPU0:router#configure
RP/0/RP0/CPU0:router(config)#router bgp 100
RP/0/RP0/CPU0:router(config-bgp)#rpki cache 172.168.35.40
RP/0/RP0/CPU0:router(config-bgp-rpki-cache)# transport ssh port 22
RP/0/RP0/CPU0:router(config-bgp-rpki-cache)#username rpki-user
RP/0/RP0/CPU0:router(config-bgp-rpki-cache)#password rpki-ssh-pass
RP/0/RP0/CPU0:router(config-bgp-rpki-cache)#preference 1
RP/0/RP0/CPU0:router(config-bgp-rpki-cache)#purge-time 30
```

rd

To configure a route distinguisher, use the **rd** command in VRF configuration mode. To disable the route distinguisher, use the **no** form of this command.

rd {*as-number* **:** *nn* | *ip-address* **:** *nn* | **auto**}
no rd {*as-number* **:** *nn* | *ip-address* **:** *nn* | **auto**}

Syntax Description

as-number:nn	• *as-number* —16-bit Autonomous system (AS) number of the route distinguisher • Range for 2-byte Autonomous system numbers (ASNs) is 1 to 65535. • Range for 4-byte Autonomous system numbers (ASNs) in asplain format is 1 to 4294967295. • Range for 4-byte Autonomous system numbers (ASNs) is asdot format is 1.0 to 65535.65535. • *nn* —32-bit number
ip-address:nn	IP address of the route distinguisher. • *ip-address* —32-bit IP address • *nn* —16-bit number
auto	Automatically assigns a unique route distinguisher.

Command Default

No default behavior or values

Command Modes

VRF configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

Use the **rd** command to make the prefix unique across multiple VRFs.

Auto assignment of route distinguishers can be done only if a router ID is assigned using the **bgp router-id** command in BGP router configuration mode. The unique router ID is used for automatic route distinguisher generation.

The following are restrictions when configuring route distinguishers:

- BGP router-id must be configured before **rd auto** can be configured
- Route distinguisher cannot be changed or removed when an IPv4 unicast address family is configured under VRF.
- BGP router-id cannot be changed or removed when **rd auto** is configured under a VRF.
- When **rd auto** is configured under a VRF, the IP address for the router distinguisher configured under another VRF must be different from that of the BGP router-id
- If a route distinguisher with same IP address as BGP router-id exists, the **rd auto** is not permitted.

Task ID

Task ID	Operations
bgp	read, write

Examples

The following example shows how to automatically assign a unique route distinguisher to VRF instance vrf-1:

```
RP/0/RP0/CPU0:router(config)# router bgp 1
RP/0/RP0/CPU0:router(config-bgp)# vrf vrf-1
RP/0/RP0/CPU0:router(config-bgp-vrf)# rd auto
```

receive-buffer-size

To set the size of the receive buffers for a Border Gateway Protocol (BGP) neighbor, use the **receive-buffer-size** command in an appropriate configuration mode. To remove the **receive-buffer-size** command from the configuration file and restore the system to its default condition in which the software uses the default size, use the **no** form of this command.

receive-buffer-size *socket-size* [*bgp-size*]
no receive-buffer-size [*socket-size*] [*bgp-size*]

Syntax Description

socket-size	Size, in bytes, of the receive-side socket buffer. Range is 512 to 131072.
bgp-size	(Optional) Size, in bytes, of the receive buffer in BGP. Range is 512 to 131072.

Command Default

socket-size : 32,768 bytes

bgp-size : 4,032 bytes

Command Modes

Neighbor configuration

VRF neighbor configuration

Neighbor group configuration

Session group configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

Use the **receive-buffer-size** command to increase the buffer size when receiving updates from a neighbor. Using larger buffers can improve convergence time because it allows the software to process a larger number of packets simultaneously. However, allocating larger buffers consumes more memory on the router.

Note Increasing the socket buffer size uses more memory only when more messages are waiting to be processed by the software. In contrast, increasing the BGP buffer size uses extra memory indefinitely.

If this command is configured for a neighbor group or session group, all neighbors using the group inherit the configuration. Values of commands configured specifically for a neighbor override inherited values.

Task ID

Task ID	Operations
bgp	read, write

Examples

The following example shows how to set the receive buffer sizes for neighbor 172.20.1.1 to be 65,536 bytes for the socket buffer and 8192 bytes for the BGP buffer:

```
RP/0/RP0/CPU0:router(config)# router bgp 1
RP/0/RP0/CPU0:router(config-bgp)# neighbor 172.20.1.1
RP/0/RP0/CPU0:router(config-bgp-nbr)# remote-as 1
RP/0/RP0/CPU0:router(config-bgp-nbr)# receive-buffer-size 65536 8192
```

redistribute (BGP)

To redistribute routes from one routing domain into Border Gateway Protocol (BGP), use the **redistribute** command in an appropriate configuration mode. To disable route redistribution, use the **no** form of this command.

Connected
redistribute connected [**metric** *metric-value*] [**route-policy** *route-policy-name*]
no redistribute connected [**metric** *metric-value*] [**route-policy** *route-policy-name*]

Intermediate System-to-Intermediate System (IS-IS)
redistribute isis *process-id* [**level** | {**1** | **1-inter-area** | **2**}] [**metric** *metric-value*] [**route-policy** *route-policy-name*]
no redistribute isis *process-id* [**level** | {**1** | **1-inter-area** | **2**}] [**metric** *metric-value*] [**route-policy** *route-policy-name*]

Open Shortest Path First (OSPF)
redistribute ospf *process-id*
no redistribute ospf *process-id*

OSPFv3
redistribute ospf *process-id*
no redistribute ospf *process-id*

Routing Information Protocol
redistribute rip [**metric** *metric-value*] [**route-policy** *route-policy-name*]
no redistribute rip [**metric** *metric-value*] [**route-policy** *route-policy-name*]

Static
redistribute static [**metric** *metric-value*] [**route-policy** *route-policy-name*]
no redistribute static [**metric** *metric-value*] [**route-policy** *route-policy-name*]

Service Layer
redistribute service-layer [**multipath**] [**route-policy** *route-policy-name*]

gRIBI
redistribute gribi [**multipath**] [**route-policy** *route-policy-name*]

Syntax Description

connected	Redistributes connected routes. Connected routes are established automatically when IP is enabled on an interface.
metric *metric-value*	(Optional) Specifies the Multi Exit Discriminator (MED) attribute used for the redistributed route. Range is 0 to 4294967295. Use a value consistent with the destination protocol. By default, the Interior Gateway Protocol (IGP) metric is assigned to the route. For connected and static routes the default metric is 0.

route-policy *route-policy-name*	(Optional) Specifies a configured routing policy to filter redistributed routes. A route policy is used to filter the importation of routes from this source routing protocol to BGP.
process-id	For the **isis** keyword, an IS-IS instance name from which routes are to be redistributed. For the **ospf** keyword, an OSPF instance name from which routes are to be redistributed. The *process-id* value takes the form of a string. A decimal number can be entered, but it is stored internally as a string.
match { **internal** | **external** [**1** | **2**] | **nssa-external** [**1** | **2**]}	(Optional) Specifies the criteria by which OSPF routes are redistributed into other routing domains. It can be one or more of the following: • **internal** —Routes that are internal to a specific autonomous system (intra- and inter-area OSPF routes). • **external** [**1 | 2**]—Routes that are external to the autonomous system, but are imported into OSPF as Type 1 or Type 2 external routes. • **nssa-external** [**1 | 2**]—Routes that are external to the autonomous system, but are imported into OSPF as Type 1 or Type 2 not-so-stubby area (NSSA) external routes. For the **external** and **nssa-external** options, if a type is not specified, then both Type 1 and Type 2 are assumed.
isis	Specifies that routes are distributed from the IS-IS protocol. Redistribution from IS-IS is allowed under IPv4 unicast, IPv4 multicast, IPv6 unicast, and IPV6 multicast address-families. Redistribution is not allowed under VPNv4 and VPNv6 address-families.
level { **1** | **1-inter-area** | **2** }	(Optional) Specifies the IS-IS level from which routes are redistributed. It can be one of the following: • **1** —Routes are redistributed from Level 1 routes. • **1-inter-area** —Routes are redistributed from Level 1 interarea routes. • **2** —Routes are redistributed from Level 2 routes.
ospf	Specifies that routes are distributed from the OSPF protocol. You must be in IPv4 unicast or multicast address family configuration mode or in VRF IPv4 address family configuration mode.
ospfv3	Specifies that routes are distributed from the OSPFv3 protocol. You must be in IPv6 unicast or multicast address family configuration mode or in VRF IPv4 address family configuration mode.
rip	Specifies that routes are distributed from RIP. You must be in IPv4 unicast or multicast address family configuration mode.
static	Redistributes IP static routes.

service-layer	The **service-layer** keyword redistributes routes programmed through the Service Layer API into BGP.
gribi	The **gribi** keyword redistributes routes programmed through gRPC RIB Interface (gRIBI) into BGP.

Command Default

Route redistribution is disabled.

For IS-IS, the default is to redistribute Level 1 and Level 2 routes.

For OSPF, the default is to redistribute internal, external, and NSSA external routes of Type 1 and Type 2.

For OSPFv3, the default is to redistribute internal, external, and NSSA external routes of Type 1 and Type 2

By default, the Interior Gateway Protocol (IGP) metric is assigned to the route. For connected and static routes the default metric is 0.

metric *metric-value:* 0

match { **internal** | **external** [**1** | **2**] | **nssa-external** [**1** | **2**]}: If no match is specified, the default is to match all routes.

Command Modes

- IPv4 address family configuration, both unicast and multicast: **connected, isis, ospf, rip,** and **static** are supported.
- IPv6 address family configuration, both unicast and multicast: **connected, isis, ospfv3,** and **static** are supported.
- VRF IPv4 address family configuration: **connected**, **ospf**, **rip**, and **static** are supported.
- VRF IPv6 address family configuration: **connected**, and **static** are supported.
- BGP address family configuration: **service-layer** and **gribi** are supported.

Command History

Release	Modification
Release 6.0	This command was introduced.
Release 26.3.1	Introduced the **service-layer** and **gribi** keywords to redistribute routes programmed through the Service Layer API and gRIBI into BGP.

Usage Guidelines

Note When redistributing routes (into BGP) using both command keywords for setting or matching of attributes and a route policy, the routes are run through the route policy first, followed by the keyword matching and setting.

Note When redistributing routes (into BGP) using **redistribute** or **redistribute connected**, configure a route-policy (RPL) to control routes.

Each instance of a protocol may be redistributed independently of the others. Changing or removing redistribution for a particular instance does not affect the redistribution capability of other protocols or other instances of the same protocol.

Networks specified using the **network** command are not affected by the **redistribute** command; that is, the routing policy specified in the **network** command takes precedence over the policy specified through the **redistribute** command.



Note Starting with Cisco IOS XR Release 26.3.1, gRPC RIB Interface (gRIBI) routes and Service Layer API routes use separate routing protocol types.

- Replace **redistribute application gribi** with **redistribute gribi** in existing configurations.
- Replace **redistribute application service-layer** with **redistribute service-layer** in existing configurations.

The **application** keyword continues to select application routes.

Task ID

Task ID	Operations
bgp	read, write

Examples

The following example shows how to redistribute IP Version 4 (IPv4) unicast OSPF routes from OSPF instance 110 into BGP:

```
RP/0/RP0/CPU0:router(config)# router bgp 109
RP/0/RP0/CPU0:router(config-bgp)# address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-bgp-af)# redistribute ospf 110
```

Examples

This example shows how to redistribute Service Layer API routes into the IPv4 unicast BGP address family.

```
Router(config)# router bgp 65000
Router(config-bgp)# address-family ipv4 unicast
Router(config-bgp-af)# redistribute service-layer route-policy SLAPI-TO-BGP
```

Examples

This example shows how to redistribute gRIBI routes into the IPv4 unicast BGP address family.

```
Router(config)# router bgp 65000
Router(config-bgp)# address-family ipv4 unicast
Router(config-bgp-af)# redistribute gribi route-policy GRIBI-TO-BGP
```

refresh-time (rpki-server)

To configure the time BGP waits in between sending periodic serial queries to the RPKI server, use the **refresh-time** command in rpki-server configuration mode. To remove the refresh-time configuration, use the **no** form of this command.

refresh-time {*time-in-seconds* | **off**}
no refresh-time {*time-in-seconds* | **off**}

Syntax Description

off	Specifies not to send serial queries periodically.
time-in-seconds	Sets the refresh-time in seconds. Range is 30 to 3600 seconds.

Command Default

Sets a default value of 600 seconds.

Command Modes

RPKI cache configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

Task ID

Task ID	Operation
bgp	read, write

This example shows how to set the refresh-time for BGP to wait in between sending periodic serial queries to the server as 30 seconds:

```
RP/0/RP0/CPU0:router#configure
RP/0/RP0/CPU0:router(config)#router bgp 100
RP/0/RP0/CPU0:router(config-bgp)#rpki server 172.168.35.40
RP/0/RP0/CPU0:router(config-bgp-rpki-cache)# transport ssh port 22
RP/0/RP0/CPU0:router(config-bgp-rpki-cache)#username rpki-user
RP/0/RP0/CPU0:router(config-bgp-rpki-cache)#password rpki-ssh-pass
RP/0/RP0/CPU0:router(config-bgp-rpki-cache)#preference 1
RP/0/RP0/CPU0:router(config-bgp-rpki-cache)#purge-time 30
RP/0/RP0/CPU0:router(config-bgp-rpki-cache)#refresh-time  30
```

remote-as (BGP)

To create a Border Gateway Protocol (BGP) neighbor and begin the exchange of routing information, use the **remote-as** command in an appropriate configuration mode. To delete the entry for the BGP neighbor, use the **no** form of this command.

remote-as *as-number*
no remote-as [*as-number*]

Syntax Description

as-number Autonomous system (AS) to which the neighbor belongs.

- Range for 2-byte Autonomous system numbers (ASNs) is 1 to 65535.
- Range for 4-byte Autonomous system numbers (ASNs) in asplain format is 1 to 4294967295.
- Range for 4-byte Autonomous system numbers (ASNs) is asdot format is 1.0 to 65535.65535.

Command Default

No BGP neighbors exist.

Command Modes

Neighbor configuration

VRF neighbor configuration

Neighbor group configuration

Session group configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

Use the **remote-as** command to create a neighbor and assign it a remote autonomous system number. A neighbor must have a remote autonomous system number before any other commands can be configured for it. Removing the remote autonomous system from a neighbor causes the neighbor to be deleted. You cannot remove the autonomous system number if the neighbor has other configuration.



Note We recommend that you use the **no neighbor** command rather than the **no remote-as** command to delete a neighbor.

A neighbor specified with a remote autonomous system number that matches the autonomous system number specified in the **router bgp** command identifies the neighbor as internal to the local autonomous system. Otherwise, the neighbor is considered external.

Configuration of the **remote-as** command for a neighbor group or session group using the **neighbor-group** command or **session-group** command causes all neighbors using the group to inherit the characteristics

configured with the command. Configuring the command directly for the neighbor overrides the value inherited from the group.

In the neighbor configuration submode, configuring use of a session group or neighbor group for which **remote-as** is configured creates a neighbor and assigns it an autonomous system number if the neighbor has not already been created.

Note Do not combine **remote-as** commands and **no use neighbor-group** commands, or **remote-as** commands and **no use session-group** commands, in the same configuration commit.

Task ID

Task ID	Operations
bgp	read, write

Examples

The following example shows how to assign autonomous system numbers on two neighbors, neighbor 10.0.0.1, (internal) and neighbor 192.168.0.1 (external), setting up a peering session that shares routing information between this router and each of these neighbors:

```
RP/0/RP0/CPU0:router(config)# router bgp 1
RP/0/RP0/CPU0:router(config-bgp)# session-group group2
RP/0/RP0/CPU0:router(config-bgp-sngrp)# remote-as 1
RP/0/RP0/CPU0:router(config-bgp-sngrp)#exit
RP/0/RP0/CPU0:router(config-bgp)# neighbor 10.0.0.1
RP/0/RP0/CPU0:router(config-bgp-nbr)#use session-group group2
```

The following example shows how to configure a session group called group2 with an autonomous system number 1. Neighbor 10.0.0.1 is created when it inherits the autonomous system number 1 from session group group2.

```
RP/0/RP0/CPU0:router(config)#router bgp 1
RP/0/RP0/CPU0:router(config-bgp)# session-group group2
RP/0/RP0/CPU0:router(config-bgp-sngrp)# remote-as 1
RP/0/RP0/CPU0:router(config-bgp-sngrp)# exit
RP/0/RP0/CPU0:router(config-bgp)# neighbor 10.0.0.1
RP/0/RP0/CPU0:router(config-bgp-nbr)# use session-group group2
```

remote-bgp-label-mode-per-prefix

To determine the most appropriate scale optimization for a specific Virtual Routing and Forwarding (VRF) for IPv4 or IPv6 prefixes, use the **remote-bgp-label-mode-per-prefix** command in global VRF configuration mode.

{ **ipv4** | **ipv6** } **remote-bgp-label-mode-per-prefix**

Syntax Description

This command has no arguments or keywords.

Command Default

None.

Command Modes

Global VRF configuration

Command History

Release	Modification
Release 7.11.1	This command was introduced.

Usage Guidelines

Use the **remote-bgp-label-mode-per-prefix** command to improve resource utilization and provide more granular control over prefix labels. This command enables the network administrators to determine whether to retain two labels in the KBP TCAM for each prefix or to utilize FEC for push labels: the choice depends on the network's specific requirements.

When you configure the **remote-bgp-label-mode-per-prefix** command under the global configuration for VRF, along with the **hw-module fib mpls two label tcam-optimized selective** command, FEC is used to push labels for all prefixes for the particular address family.

Task ID	Operation
bgp	read, write

Example

The following example shows how to configure scale optimization for a specific VRF for IPv4 prefix:

```
Router(config)# vrf vrf1
Router(config-vrf)# ipv4 remote-bgp-label-mode-per-prefix

Router(config-vrf)# address-family ipv4 unicast
Router(config-vrf-af)# import route-target 1:1
Router(config-vrf-af)# export route-target 1:1
Router(config-vrf-af)# exit
```

The following example shows how to configure scale optimization for a specific VRF for IPv6 prefix:

```
Router(config)# vrf vrf1
Router(config-vrf)# ipv6 remote-bgp-label-mode-per-prefix

Router(config-vrf)# address-family ipv4 unicast
Router(config-vrf-af)# import route-target 1:1
```

```
Router(config-vrf-af)# export route-target 1:1
Router(config-vrf-af)# exit
```

remove-private-as

To remove private autonomous system numbers from autonomous system paths when generating updates to external neighbors, use the **remove-private-as** command in an appropriate configuration mode. To place the router in the default state in which it does not remove private autonomous system numbers, use the **no** form of this command.

remove-private-as [**inheritance-disable**] [**entire-aspath**]
no remove-private-as [**inheritance-disable**] [**entire-aspath**]

Syntax Description

inheritance-disable	(Optional) Permits the feature to be disabled from a neighbor group or address family group instead of being inherited.
entire-aspath	(Optional) Removes the entire private autonomous system numbers from an autonomous system path only if all ASes in the path are private.

Command Default

When this command is not specified in the appropriate configuration mode, private autonomous system numbers are not removed from updates sent to external neighbors.

Command Modes

IPv4 address family group configuration

IPv6 address family group configuration

IPv4 neighbor address family configuration

IPv4 neighbor group address family configuration

IPv6 neighbor group address family configuration

VRF IPv4 neighbor address family configuration

VRF IPv4 neighbor address family configuration

VPNv6 address family group configuration

VPNv6 neighbor address family configuration

VPNv6 neighbor address family configuration

VRF IPv6 neighbor address family configuration

VPNv6 neighbor group address family configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

This feature is available for external BGP (eBGP) neighbors only.

When an update is passed to the external neighbor, the software drops any private autonomous system numbers, no matter it is leading or in the middle of the AS_SEQUENCE.

If this command is used in a BGP confederation, the element following the confederation portion of the autonomous system path, if a sequence, is considered the leading sequence.

The private autonomous system values range from 64512 to 65535.

If this command is configured for a neighbor group or address family group, all neighbors using the group inherit the configuration. Values of commands configured specifically for a neighbor override inherited values.

Use the **entire-aspath** to removes the entire private autonomous system numbers from an autonomous system path only if all ASes in the path are private.

Task ID

Task ID	Operations
bgp	read, write

Examples

The following example shows a configuration that removes the private autonomous system number from the IP Version 4 (IPv4) unicast updates sent to 172.20.1.1:

```
RP/0/RP0/CPU0:router(config)# router bgp 140
RP/0/RP0/CPU0:router(config-bgp)# neighbor 172.20.1.1
RP/0/RP0/CPU0:router(config-bgp-nbr)# remote-as 1
RP/0/RP0/CPU0:router(config-bgp-nbr)# address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-bgp-nbr-af)# remove-private-as
```

The following example shows how to disable the remove private autonomous system number feature for neighbor 172.20.1.1, preventing this feature from being automatically inherited from address family group group1:

```
RP/0/RP0/CPU0:router(config)# router bgp 140
RP/0/RP0/CPU0:router(config-bgp)# af-group group1 address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-bgp-afgrp)# remove-private-as
RP/0/RP0/CPU0:router(config-bgp-afgrp)# exit
RP/0/RP0/CPU0:router(config-bgp)# neighbor 172.20.1.1
RP/0/RP0/CPU0:router(config-bgp-nbr# remote-as 1
RP/0/RP0/CPU0:router(config-bgp-nbr)# address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-bgp-nbr-af)# use af-group group1
RP/0/RP0/CPU0:router(config-bgp-nbr-af)# remove-private-as inheritance-disable
```

replace as-path all

To replace BGP AS path with custom values, use the **replace as-path all** command in route-policy configuration mode.

replace as-path all { **none auto** | *'x'* } [*n*] [*parameter*]

Syntax Description

none	Replaces AS path with null or empty value.
auto	Replaces AS path based on the route policy. For inbound route policy, replaces AS path with AS path of neighbor from where the prefix is received. For outbound route policy, replaces AS path with the configured local AS path.
'x'	Replaces AS path with specified value, where *'x'* is a single AS number or a sequence of AS numbers separated by space.
n	(Optional) Repeats replacing the AS path for specified number of times. The range is from 2 to 64. This option is supported only for the **auto** and *'x'* parameters.
parameter	(Optional) Parameter name used along with repeat option. The parameter name must be preceded with a "$". You can attach the route policy configured with the parameter to a neighbor and specify the number of times the AS path replacement should be repeated.

Command Default

No default behavior or values

Command Modes

Route-policy configuration

Command History

Release	Modification
Release 7.5.2	This command was introduced.

Usage Guidelines

Use the **replace as-path all** command to replace the BGP AS path with specified custom values.

When you remove the route policy from a neighbor, the associated **replace as-path all** configuration is also removed.

Task ID

Task ID	Operations
route-policy	read, write

The following example shows how to use the **replace as-path all** command to replace the AS path with **auto**:

```
Router(config)#route-policy aspath-auto
Router(config-rpl)# replace as-path all auto
Router(config-rpl)# end-policy
Router(config)# router bgp 65530
```

```
Router(config-bgp)# neighbor 111.0.0.1
Router(config-bgp-nbr)# address-family ipv4 unicast
Router(config-bgp-nbr-af)# route-policy aspath-auto out
```

response-time (rpki-server)

To configure the time BGP waits for a response from the RPKI cache-server after sending a serial or reset query, use the **response-time** command in rpki-server configuration mode. To remove the response-time configuration, use the **no** form of this command.

response-time {*time-in-seconds* | **off**}
no response-time {*time-in-seconds* | **off**}

Syntax Description

off	Specifies to wait indefinitely for a response from the RPKI cache.
time-in-seconds	Specifies the response-time in seconds. Range is 30 to 3600 seconds.

Command Default

Sets a default value of 30 seconds.

Command Modes

RPKI server configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

Task ID

Task ID	Operation
bgp	read, write

This example shows how to set the time for BGP to wait for a response from the RPKI server as 30 seconds, after sending a serial or reset query:

```
RP/0/RP0/CPU0:router#configure
RP/0/RP0/CPU0:router(config)#router bgp 100
RP/0/RP0/CPU0:router(config-bgp)#rpki server 72.168.35.40
RP/0/RP0/CPU0:router(config-bgp-rpki-cache)# transport ssh port 22
RP/0/RP0/CPU0:router(config-bgp-rpki-cache)#username rpki-user
RP/0/RP0/CPU0:router(config-bgp-rpki-cache)#password rpki-ssh-pass
RP/0/RP0/CPU0:router(config-bgp-rpki-cache)#preference 1
RP/0/RP0/CPU0:router(config-bgp-rpki-cache)#purge-time 30
RP/0/RP0/CPU0:router(config-bgp-rpki-cache)#refresh-time 30
RP/0/RP0/CPU0:router(config-bgp-rpki-cache)#response-time 30
```

retain local-label

To retain the local label until the network is converged, use the **retain local-label** command in an appropriate address family configuration mode. To disable the retaining of the local label, use the **no** form of this command.

retain local-label *minutes*
no retain local-label

Syntax Description

minutes	Local retention time in minutes. The range is 3 to 60 minutes. The default retention time is 5 minutes.

Command Default

minutes : 5

Command Modes

L2VPN address family configuration

VPNv4 address family configuration

VPNv6 address family configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

When configuring the **label mode per-nexthop-received-label** command at the VPN level, you must explicitly configure the local retention time for the **retain local-label** command.

Task ID

Task ID	Operations
bgp	read, write

Examples

The following example shows how to enable local label retention for 5 minutes:

```
Router# configure
Router(config)# router bgp 100
Router(config-bgp)# address-family vpnv4 unicast
Router(config-bgp-af)# retain local-label 5
```

retain route-target

To accept received updates with specified route targets, use the **retain route-target** command in an appropriate configuration mode. To disable the retaining of routes tagged with specified route targets, use the **no** form of this command.

retain route-target {**all** | **route-policy** *route-policy-name*}
no retain route-target [**all** | **route-policy** *route-policy-name*]

Syntax Description

all	Accepts received updates containing at least one route target.
route-policy *router-policy-name*	Accepts received updates accepted by a specified route filter policy.

Command Default

The default is to accept all route targets.

Command Modes

VPNv4 address family configuration

VPNv6 address family configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

Use the **retain route-target** command to configure a route reflector (RR) to retain routes tagged with specific route targets (RT).

The **retain route-target** is a required command for Inter-AS option B ASBR. By default, an Inter-AS option B ASBR needs the **retain route-target** configured to get VPNv4 BGP table from PE routers, either with the **all** or with the **route-policy** option.

A provider edge (PE) router is not required to hold all VPNv4 routes. The PE router holds only routes that match the import RT of the VPNs configured on it, but a RR must retain all VPNv4 routes because it may peer with PE routers and different PEs may require different RT-tagged VPNv4 routes. Configuring an RR to hold only routes that have a defined set of RT communities and configuring some of these RRs to service a different set of VPNs provides scalability to the RRs. A PE can be configured to peer with all RRs that service the VPN routing and forwarding (VRF) instances configured on the PE. When a new VRF is configured with an RT for which the PE does not already hold routes, the PE issues route refresh requests to the RRs and gets the relevant VPN routes.

The **route-policy** *route-policy-name* keyword and argument takes the policy name that lists the extended communities that a path should have for the RR to retain the path.

Task ID

Task ID	Operations
bgp	read, write

Examples

The following example shows how to configure RR to retain all routes with the route filter policy ft-policy-A:

```
RP/0/RP0/CPU0:router(config)# router bgp 140
RP/0/RP0/CPU0:router(config-bgp)# address-family vpnv4 unicast
RP/0/RP0/CPU0:router(config-bgp-af)# retain route-target route-filter ft-policy-A
```

route-policy (BGP)

To apply a routing policy to updates advertised to or received from a Border Gateway Protocol (BGP) neighbor, use the **route-policy** command in an appropriate configuration mode. To disable applying routing policy to updates, use the **no** form of this command.

route-policy *route-policy-name* [*parameter1, parameter2, . . . , parametern*] {**in** | **out**}
no route-policy *route-policy-name* [*parameter1, parameter2, . . . , parametern*] {**in** | **out**}

Syntax Description

route-policy-name	Name of route policy. Up to 16 parameters can follow the route-policy-name, enclosed in brackets ([ ]).
in	Applies policy to inbound routes.
out	Applies policy to outbound routes.

Command Default

No policy is applied.

Command Modes

IPv4 address family group configuration

IPv6 address family group configuration

IPv4 neighbor address family configuration

IPv4 neighbor group address family configuration

IPv6 neighbor group address family configuration

VPNv4 address family group configuration

VPNv4 neighbor address family configuration

VRF IPv4 neighbor address family configuration

VPNv4 neighbor group address family configuration

VPNv6 address family group configuration

VRF IPv6 neighbor address family configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

Use the **route-policy** command to specify a routing policy for an inbound or outbound route. The policy can be used to filter routes or modify route attributes. The **route-policy** command is used to define a policy.

Note Configuring a large number of uniquely named outbound neighbor policies can adversely affect performance. This is true even if the uniquely named route policies are functionally identical. The user is discouraged from configuring multiple functionally identical route policies for use with this command. For example, if Policy A and Policy B are identical but named for different neighbors, the two policies should be configured as a single policy.

If the **route-policy** command is configured for a neighbor group or neighbor address family group, all neighbors using the group inherit the configuration. Values of commands configured specifically for a neighbor override inherited values.

Task ID

Task ID	Operations
bgp	read, write

Examples

The following example shows how to apply the In-Ipv4 policy to inbound IP Version 4 (IPv4) unicast routes from neighbor 172.20.1.1:

```
RP/0/RP0/CPU0:router(config)# router bgp 1
RP/0/RP0/CPU0:router(config-bgp)# neighbor 172.20.1.1
RP/0/RP0/CPU0:router(config-bgp-nbr)# remote-as 1
RP/0/RP0/CPU0:router(config-bgp-nbr)# address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-bgp-nbr-af)# route-policy In-Ipv4 in
```

router bgp

To configure the Border Gateway Protocol (BGP) routing process, use the **router bgp** command in XR Config mode. To remove all BGP configurations and terminate the BGP routing process, use the **no** form of this command.

router bgp *as-number* [**instance** *instance-name*]
no router bgp [*as-number*]

Syntax Description

as-number	Number that identifies the autonomous system (AS) in which the router resides. • Range for 2-byte Autonomous system numbers (ASNs) is 1 to 65535. • Range for 4-byte Autonomous system numbers (ASNs) in asplain format is 1 to 4294967295. • Range for 4-byte Autonomous system numbers (ASNs) is asdot format is 1.0 to 65535.65535.
instance *instance-name*	Specifies an instance and instance name. The maximum length for the instance name is 32 characters. The **router bgp instance** *instance-name* command replaced the **distributed speaker** command.

Command Default

No BGP routing process is enabled.

Command Modes

XR Config mode

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

Use the **router bgp** command to set up a distributed routing core that automatically guarantees the loop-free exchange of routing information between autonomous systems.

Task ID

Task ID	Operations
bgp	read, write
rib	read, write

Examples

The following example shows how to configure a BGP process for autonomous system 120:

```
RP/0/RP0/CPU0:router(config)# router bgp 120
```

rpki server

To enter resource public key infrastructure (RPKI) cache-server (rpki-sever) configuration mode and enable rpki parameters configuration, use the **rpki server** command in Router BGP configuration mode. To remove the rpki-server configuration mode and delink cache-server from the cache list, use the **no** form of this command.

rpki server {*host-nameip-address*}
no rpki server {*host-nameip-address*}

Syntax Description

host-name	Host name of the RPKI cache database.
ip-address	IP Address of the RPKI cache databse.

Command Default

RPKI server configuration is disabled.

Command Modes

Router BGP configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

Task ID

Task ID	Operation
bgp	read, write

This example shows how to configure an rpki cache-server database and enter rpki-server configuration mode:

```
RP/0/RP0/CPU0:router#configure
RP/0/RP0/CPU0:router(config)#router bgp 100
RP/0/RP0/CPU0:router(config-bgp)#rpki server 172.168.35.40
RP/0/RP0/CPU0:router(config-bgp-rpki-cache)#
```

selective-vrf-download disable

To disable selective VRF download (SVD) on a line card to enable download all prefixes and labels to the line card, use the **selective-vrf-download disable** command in XR Config mode. To enable the SVD, use the **no** form of this command.

selective-vrf-download disable
no selective-vrf-download disable

Syntax Description

This command has no keywords or arguments.

Command Default

SVD is enabled.

Command Modes

XR Config mode

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

You must failover the active RP or reload the router after disabling SVD for the configuration change to get activated.

Task ID

Task ID	Operation
ip-services	read, write

This example shows how to disable selective vrf download:

```
RP/0/RP0/CPU0:router#configure
RP/0/RP0/CPU0:router(config)#selective-vrf-download disable
```

send-buffer-size

To set the size of the send buffers for a Border Gateway Protocol (BGP) neighbor, use the **send-buffer-size** command in an appropriate configuration mode. To set the size of the send buffers to the default values, use the **no** form of this command.

send-buffer-size *socket-size* [*bgp-size*]
no send-buffer-size [*socket-size*] [*bgp-size*]

Syntax Description

socket-size	Size, in bytes, of the send-side socket buffer. Range is 4096 to 131072.
bgp-size	(Optional) Size, in bytes, of the BGP process send buffer. Range is 4096 to 131072.

Command Default

socket-size : 10240 bytes

bgp-size : 4096 bytes

Use the **socket send-buffer-size** command to change the defaults.

Command Modes

Neighbor configuration

VRF neighbor configuration

Neighbor group configuration

Session group configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

Use the **send-buffer-size** command to increase the buffer size employed when sending updates to a neighbor. Using larger buffers can improve convergence time because the software can process more packets simultaneously. However, allocating larger buffers uses more memory on the router.

Note Increasing the socket buffer size uses more memory only when more messages are waiting to be processed by the software. In contrast, increasing the BGP buffer size uses more memory indefinitely.

If this command is configured for a neighbor group or session group, all neighbors using the group inherit the configuration. Values of commands configured specifically for a neighbor override inherited values.

Task ID

Task ID	Operations
bgp	read, write

Examples

The following example shows how to set the send buffer sizes for neighbor 172.20.1.1 to be 8192 bytes for both the socket buffer and the BGP buffer:

```
RP/0/RP0/CPU0:router(config)# router bgp 1
RP/0/RP0/CPU0:router(config-bgp)# neighbor 172.20.1.1
RP/0/RP0/CPU0:router(config-bgp-nbr)# remote-as 1
RP/0/RP0/CPU0:router(config-bgp-nbr)# send-buffer-size 8192 8192
```

send-community-ebgp

To specify that community or large community attributes should be sent to an external Border Gateway Protocol (eBGP) neighbor, use the **send-community-ebgp** command in an appropriate configuration mode. To disable sending community or large community attributes to an eBGP neighbor, use the **no** form of this command.

send-community-ebgp [**inheritance-disable**]
no send-community-ebgp [**inheritance-disable**]

Syntax Description

inheritance-disable	(Optional) Allows configuration inherited from a neighbor group or address family group to be overridden.

Command Default

Community (COMM) attributes are NOT sent to eBGP peers (including PE-CE peers).

Command Modes

IPv4 address family group configuration

IPv6 address family group configuration

IPv4 neighbor address family configuration

IPv4 neighbor group address family configuration

IPv6 neighbor group address family configuration

VRF IPv4 neighbor address family configuration

VPNv4 neighbor address family configuration

VRF IPv6 neighbor address family configuration

VPNv6 neighbor address family configuration

Command History

Release	Modification
Release 6.0	This command was introduced.
Release 6.3.1	This command supports large communities.

Usage Guidelines

Use the **send-community-ebgp** command to control whether community or large community attributes are sent to eBGP neighbors. This command cannot be configured for iBGP neighbors as community or large community attributes are always sent to iBGP neighbors.

When IOS XR BGP updates community or large community attributes for eBGP VPN peers (VPNv4 or VPNv6), there is no need to configure the **send-community-ebgp** command separately. The community or large community attributes are updated by default.

If this command is configured for a neighbor group or address family group, all neighbors using the group inherit the configuration. Configuring the command specifically for a neighbor overrides inherited values.

Task ID

Task ID	Operations
bgp	read, write

Examples

The following example shows how to disable the router that sends community attributes to neighbor 172.20.1.1 for IP Version 4 (IPv4) multicast routes:

```
RP/0/RP0/CPU0:router(config)#router bgp 140
RP/0/RP0/CPU0:router(config-bgp)# neighbor 172.20.1.1
RP/0/RP0/CPU0:router(config-bgp-nbr)# remote-as 1
RP/0/RP0/CPU0:router(config-bgp-nbr)# address-family ipv4 multicast
RP/0/RP0/CPU0:router(config-bgp-nbr-af)# send-community-ebgp
```

The following example shows how to disable the delivery of community attributes to neighbor 172.20.1.1, preventing this feature from being inherited from address family group group1:

```
RP/0/RP0/CPU0:router(config)#router bgp 140
RP/0/RP0/CPU0:router(config-bgp)# af-group group1 address-family ipv4 multicast
RP/0/RP0/CPU0:router(config-bgp-afgrp)# send-community-ebgp
RP/0/RP0/CPU0:router(config-bgp-afgrp)# exit
RP/0/RP0/CPU0:router(config-bgp)# neighbor 172.20.1.1
RP/0/RP0/CPU0:router(config-bgp-nbr)# remote-as 1
RP/0/RP0/CPU0:router(config-bgp-nbr)# address-family ipv4 multicast
RP/0/RP0/CPU0:router(config-bgp-nbr-af)# use af-group group1
RP/0/RP0/CPU0:router(config-bgp-nbr-af)# send-community-ebgp inheritance-disable
```

send-extended-community-ebgp

To specify that extended community attributes should be sent to external Border Gateway Protocol (eBGP) neighbors, use the **send-extended-community-ebgp** command in an appropriate configuration mode. To disable sending extended community attributes to eBGP neighbors, use the **no** form of this command.

send-extended-community-ebgp [**inheritance-disable**]
no send-extended-community-ebgp [**inheritance-disable**]

Syntax Description

inheritance-disable	(Optional) Allows configurations inherited from a neighbor group or address family group to be overridden.

Command Default

Extended community (EXTCOMM) attributes are NOT sent to eBGP peers (including PE-CE peers).

Command Modes

IPv4 address family group configuration

IPv6 address family group configuration

IPv4 neighbor address family configuration

IPv4 neighbor group address family configuration

IPv6 neighbor group address family configuration

VRF IPv4 neighbor address family configuration

VPNv4 neighbor address family configuration

VRF IPv6 neighbor address family configuration

VPNv6 neighbor address family configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

Use the **send-extended-community-ebgp** command to control whether extended community attributes are sent to eBGP neighbors. This command cannot be used for iBGP neighbors as extended community attributes are always sent to iBGP neighbors.

When IOS XR BGP updates community attributes for eBGP VPN peers (VPNv4 or VPNv6), there is no need to configure the **send-extended-community-ebgp** command separately. The community attributes are updated by default.

If this command is configured for a neighbor group or neighbor address family group, all neighbors using the group inherit the configuration. Values of commands configured specifically for a neighbor override inherited values.

Task ID

Task ID	Operations
bgp	read, write

Examples

The following example shows how to configure the router to send extended community attributes to neighbor 172.20.1.1 for IP Version 4 (IPv4) multicast routes:

```
RP/0/RP0/CPU0:router(config)# router bgp 140
RP/0/RP0/CPU0:router(config-bgp)# neighbor 172.20.1.1
RP/0/RP0/CPU0:router(config-bgp-nbr)# remote-as 1
RP/0/RP0/CPU0:router(config-bgp-nbr)# address-family ipv4 multicast
RP/0/RP0/CPU0:router(config-bgp-nbr-af)# send-extended-community-ebgp
```

The following example shows how to disable the delivery of extended community attributes to neighbor 172.20.1.1, preventing this feature from being automatically inherited from address family group group1:

```
RP/0/RP0/CPU0:router(config)# router bgp 140
RP/0/RP0/CPU0:router(config-bgp)# af-group group1 address-family ipv4 multicast
RP/0/RP0/CPU0:router(config-bgp-afgrp)# send-extended-community-ebgp
RP/0/RP0/CPU0:router(config-bgp-afgrp)# exit
RP/0/RP0/CPU0:router(config-bgp)# neighbor 172.20.1.1
RP/0/RP0/CPU0:router(config-bgp-nbr)# remote-as 1
RP/0/RP0/CPU0:router(config-bgp-nbr)# address-family ipv4 multicast
RP/0/RP0/CPU0:router(config-bgp-nbr-af)# use af-group group1
RP/0/RP0/CPU0:router(config-bgp-nbr-af)# send-extended-community-ebgp inheritance-disable
```

session-group

To create a session group and enter session group configuration mode, use the **session-group** command inXR Config mode. To remove a session group and delete all configurations associated with it, use the **no** form of this command.

session-group *name*
no session-group *name*

Syntax Description

name	Name of the session group.

Command Default

No session groups are created.

Command Modes

XR Config mode

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

Use the **session-group** command to create a session group from which neighbors can inherit configuration that is address family-independent. That is, session groups cannot have address family-specific configuration. This command enters the session group configuration mode in which configuration for a session group is entered.

Many commands can be configured in both session group configuration mode and neighbor configuration mode.

Use of session groups saves time and reduces the router configuration size. Because the configuration of a session group can be inherited by any number of neighbors, use of the group can eliminate the need to copy long or complex configurations on each of a large number of neighbors. A neighbor can inherit all configuration from a session group simply by configuring the **use** command. Specific inherited session group configuration commands can be overridden for a specific neighbor by explicitly configuring the command for the specific neighbor.

The **no** form of this command causes all of the configuration for the session group to be removed. You cannot use the **no** form of this command if removing the group would leave one or more neighbors without a configured remote autonomous system number.

Task ID

Task ID	Operations
bgp	read, write

Examples

The following example shows a session group called group1 that is used by two neighbors, 10.0.0.1 and 10.0.0.2. Because group1 is a session group, it contains only address family-independent configuration. And because group1 is used by neighbors 10.0.0.1 and 10.0.0.2, they inherit the configuration of the group.

```
RP/0/RSP0RP0/CPU0:router(config)# router bgp 1
RP/0/RSP0RP0/CPU0:router(config-bgp)# session-group group1
RP/0/RSP0RP0/CPU0:router(config-bgp-sngrp)# remote-as 1
RP/0/RSP0RP0/CPU0:router(config-bgp-sngrp)# advertisement-interval 2
RP/0/RSP0RP0/CPU0:router(config-bgp-sngrp)# exit
RP/0/RSP0RP0/CPU0:router(config-bgp)# neighbor 10.0.0.1
RP/0/RSP0RP0/CPU0:router(config-bgp-nbr)# use session-group group1
RP/0/RSP0RP0/CPU0:router(config-bgp-nbr)# exit
RP/0/RSP0RP0/CPU0:router(config-bgp)# neighbor 10.0.0.2
RP/0/RSP0RP0/CPU0:router(config-bgp-nbr)# use session-group group1
```

The following example shows a session group called group1 used by two neighbors, 10.0.0.1 and 10.0.0.2. Because group1 is a session group, it contains only address family-independent configuration. And because group1 is used by neighbors 10.0.0.1 and 10.0.0.2, they inherit the configuration of the group. However, the **password password1** configuration from group1 is overridden for neighbor 10.0.0.2, using the **password-disable** command in the neighbor 10.0.0.2 configuration submode.

```
RP/0/RP0/CPU0:router(config)# router bgp 1
RP/0/RP0/CPU0:router(config-bgp)# session-group group1
RP/0/RP0/CPU0:router(config-bgp-sngrp)# remote-as 1
RP/0/RP0/CPU0:router(config-bgp-sngrp)# advertisement-interval 2
RP/0/RP0/CPU0:router(config-bgp-sngrp)# password password1
RP/0/RP0/CPU0:router(config-bgp-sngrp)# exit
RP/0/RP0/CPU0:router(config-bgp)# neighbor 10.0.0.1
RP/0/RP0/CPU0:router(config-bgp-nbr)# use session-group group1
RP/0/RP0/CPU0:router(config-bgp-nbr)# exit
RP/0/RP0/CPU0:router(config-bgp)# neighbor 10.0.0.2
RP/0/RP0/CPU0:router(config-bgp-nbr)# use session-group group1
RP/0/RP0/CPU0:router(config-bgp-nbr)# password-disable
```

session-open-mode

To establish a Border Gateway Protocol (BGP) session with a specific TCP open mode, use the **session-open-mode** command in an appropriate configuration mode. To restore the default state, use the **no** form of this command.

session-open-mode {**active-only** | **both** | **passive-only**}
no session-open-mode [**active-only** | **both** | **passive-only**]

Syntax Description

active-only	Ensures that the BGP session can be established only when the request is initiated by the local end (active-open request) and all passive-open requests (from the other end) are rejected by the local BGP.
both	Allows BGP sessions to be established from both incoming or outgoing TCP connection requests, with one being rejected in the event of a request collision.
passive-only	Ensures that the local BGP does not initiate any TCP open requests and the session can be established only when the request comes from the remote end.

Command Default

The default is **both** .

Command Modes

Neighbor configuration

VRF neighbor configuration

Neighbor group configuration

Session group configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

BGP, by default, tries to initiate an active TCP connection whenever a new neighbor is configured. A remote neighbor may also initiate the TCP connection before the local BGP can initiate the connection. This initiation of a TCP connection by a remote neighbor is considered a passive-open request and it is accepted by the local BGP. This default behavior can be modified using the **session-open-mode** command.



Note The BGP connection is not opened and, as a result the BGP session, is not established if both the peering neighbors use the same nondefault TCP session open mode—active-only or passive-only. If both ends are configured with active-only, each neighbor rejects the TCP open request from the other end. One neighbor must be configured as passive-only or both. Similarly, if both neighbors are configured with passive-only, neither neighbor initiates the TCP open request and the BGP session is not established. Again, one neighbor must be configured as active-only or both. There is one exception. A connection open request from a neighbor that is configured with the TCP session open mode to be passive-only is processed to detect whether there is a connection collision before the request is rejected. This exception enables the local BGP to reset the session if the remote neighbor goes down and it is not detected by the local router.

Use the **session-open-mode** command when it may be necessary to preconfigure a neighbor that does not exist. Ensure that BGP does not spend any time actively trying to set up a TCP session with the neighbor. A BGP session does not come up between two neighbors, both of which configure the same nondefault value (**active-only** or **passive-only** keyword) for this command.

Task ID

Task ID	Operations
bgp	read, write

Examples

The following example shows how to enable a BGP session on router bgp 1:

```
RP/0/RP0/CPU0:router(config)# router bgp 1
RP/0/RP0/CPU0:router(config-bgp)# neighbor 45.67.89.01
RP/0/RP0/CPU0:router(config-bgp-nbr)# session-open-mode active-only
```

set flow-tag

To set the flow-tag value for the PBR BGP, use the **set flow-tag** command in route-policy configuration mode.

set flow-tag {*numberparameter*}

Syntax Description

number	Flow-tag value. Range is from 1 to 63.
parameter	Parameter name. The parameter name must be preceded with a "$."

Command Default

No default behavior or values

Command Modes

Route-policy configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

Use the **set flow-tag** command to set the flow-tag to classify packets.

This command is supported at the BGP table-policy attachpoint. Prefixes are marked for subsequent processing in the forwarding plane. After flow-tag propagation through Border Gateway Protocol (BGP), flow-tag is enabled on an interface, corresponding traffic shaping and policing is completed using packet classification based on the flow-tag value.

Task ID

Task ID	Operations
route-policy	read, write

Examples

This example shows how to use **set flow-tag** command:

```
RP/0/RP0/CPU0:router(config)# route-policy policy_1
RP/0/RP0/CPU0:router(config-rpl)# set flow-tag 12
RP/0/RP0/CPU0:router(config-rpl)# end-policy
```

set path-color external-reach

To set the BGP external-reach value for routing information base (RIB), use the **set path-color external-reach** command in route-policy configuration mode.

set path-color external-reach

Syntax Description This command has no keywords or arguments.

Command Default No default behavior or values

Command Modes Route-policy configuration

Command History

Release	Modification
Release 6.1.2	This command was introduced.

Usage Guidelines None.

Task ID

Task ID	Operations
route-policy	read, write

This example shows how to use **set path-color external-reach** command:

```
RP/0/RP0/CPU0:router(config)#route-policy HILO_FIB
RP/0/RP0/CPU0:router(config-rpl)#if destination in (150.0.0.0/8 le 24) then
RP/0/RP0/CPU0:router(config-rpl-if)#set path-color external-reach
RP/0/RP0/CPU0:router(config-rpl-if)#pass
RP/0/RP0/CPU0:router(config-rpl-if)#else
RP/0/RP0/CPU0:router(config-rpl-else)#pass
RP/0/RP0/CPU0:router(config-rpl-else)#endif
RP/0/RP0/CPU0:router(config-rpl)#end-policy
RP/0/RP0/CPU0:router(config)#commit
```

show bgp

To display entries in the Border Gateway Protocol (BGP) routing table, use the **show bgp** command in XR EXEC mode.

show bgp [**ipv4** { **unicast** | **multicast** | **labeled-unicast** | **all** | **tunnel** | **mdt** | **flowspec** | **rt-filter** } | **ipv6** {**unicast** | **multicast** | **all** | **labeled-unicast** | | **flowspec**} | **all** {**unicast** | **multicast** | **all** | **labeled-unicast** | **mdt** | **tunnel**} | **vpnv4** { **flowspec** | **multicast** | **unicast** | **rt-set**} [**rd** *rd-address*] | **vrf** {*vrf-name* | **all**} [**ipv4** {**unicast** | **labeled-unicast**} | **ipv6** **{unicast | flowspec}**] | **vpnv6** { **flowspec** | **unicast** | **rt-set**} | [**instance**] | [**instances**] | **flowspec**] [*ip-address* [{*mask* | */prefix-length*} [**longer-prefixes** | **unknown-attributes** | **bestpath-compare**]]] [**standby**] [**detail**] **rt-set**

Syntax Description

ipv4	(Optional) Specifies IP Version 4 address prefixes.
unicast	(Optional) Specifies unicast address prefixes.
multicast	(Optional) Specifies multicast address prefixes.
labeled-unicast	(Optional) Specifies labeled unicast address prefixes.
all	(Optional) For subaddress families, specifies prefixes for all subaddress families.
tunnel	(Optional) Specifies tunnel address prefixes.
ipv6	(Optional) Specifies IP Version 6 address prefixes.
all	(Optional) For address family, specifies prefixes for all address families.
vpnv4 unicast	(Optional) Specifies VPNv4 unicast address families.
instances	(Optional) Displays information of all BGP instances.
vrf	(Optional) Specifies VPN routing and forwarding (VRF) instance.
vrf	(Optional) Specifies VPN routing and forwarding (VRF) instance.
vrf-name	(Optional) Name of a VRF.
all	(Optional) For VRF, specifies all VRFs.
ipv4 { **unicast** | **labeled-unicast** }	(Optional) For VRF, specifies IPv4 unicast or labeled-unicast address families.
ipv6 unicast	(Optional) For VRF, specifies IPv6 unicast address families.
ip-address	(Optional) Network address, entered to display a particular network in the BGP routing table. If the network address is omitted, then all networks in the BGP routing table are displayed. If the network mask and prefix length is omitted, then the software displays the longest matching prefix for the network address.
mask	(Optional) Network mask of the BGP route to match.

/ prefix-length	(Optional) Prefix length of the BGP route to match. A slash (/) must precede the decimal value.
longer-prefixes	(Optional) Displays a route with the specified prefix length and more-specific routes if available. The **longer-prefixes** keyword is available when the *ip-address* and *mask* or */prefix-length* arguments are specified.
unknown-attributes	(Optional) Includes unknown, transitive attributes. The **unknown-attributes** keyword is available when the *ip-address* and *mask* or */prefix-length* arguments are specified.
bestpath-compare	(Optional) Displays route and best-path comparison information. The **bestpath-compare** keyword is available when the *ip-address* and *mask* or */prefix-length* arguments are specified.
rt-set	Displays all RT-sets for a given address-family.
flowspec	Displays flowspec configuration information.
vpnv4 multicast	Displays VPNv4 multicast prefixes.
ipv4 rt-filter	Verifies the the RT-constraint default prefix state for specified neighbor.

Command Default

If no address family or subaddress family is specified, the default address family and subaddress family specified using the **set default-afi** and **set default-safi** commands are used.

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 26.3.1	A new keyword, **rt-filter**, was introduced.
Release 6.3.1	The command output was modified to include large community in the list of attributes.
Release 6.0	This command was introduced.

Usage Guidelines

Note The **set default-afi** command is used to specify the default address family for the sessions and the **set default-safi** command is used to specify the default subaddress family for the session. See the for detailed information and syntax for the **set default-afi** and **set default-safi** commands. If you do not specify a default address family, the default address family is IPv4. If you do not specify a default subaddress family, the default subaddress family is unicast.

BGP contains a separate routing table for each address family and subaddress family combination that has been configured. The address family and subaddress family options specify the routing table to be examined. If the **all** keyword is specified for an address family or a subaddress family, each matching routing table is examined in turn.

Note Running the **show bgp** command immediately after configuring a large and complex route policy may result in timeout of the system database shown through an error message (`SYSDB-SYSDB-6-TIMEOUT_EDM`). It is recommended, that the show command be run, after the new route policy takes effect.

Use the **show bgp** *ip-address* { *mask* | */ prefix-length* } command to display detailed information for a specific route. If the mask and prefix length are omitted, the details of the longest matching prefix for the IP address are displayed.

Use the **show bgp** command to display all routes in the specified BGP routing table. Use the **show bgp** *ip-address* { *mask* | */ prefix-length* } **longer-prefixes** command to display those routes more specific than a particular prefix.

Use the **unknown-attributes** keyword to display details of any transitive attributes associated with a route that are not understood by the local system.

Use the **show bgp** *ip-address/prefix-length* **detail** command to display details of the specified prefix.

The **show bgp ipv4 rt-filter** command is supported only on NCS 540 routers.

Task ID

Task ID	Operations
bgp	read

Examples

The following is sample output from the **show bgp** command in XR EXEC mode with the BGP Persistence or long lived graceful restart (LLGR) status:

```
RP/0/RP0/CPU0:router# show bgp vpnv4 uni rd 2:1 3.0.0.0/24
[KBGP routing table entry for 3.0.0.0/24, Route Distinguisher: 2:1
Versions:
  Process           bRIB/RIB  SendTblVer
  Speaker             350584      350584
    Local Label: 16010
Last Modified: Jun 23 06:22:12.821 for 00:03:27
Paths: (1 available, best #1)
  Not advertised to any peer
  Path #1: Received by speaker 0
  Not advertised to any peer
  6913, (Received from a RR-client), (long-lived stale)
    4.4.4.4 (metric 3) from 3.3.3.3 (4.4.4.4)
      Received Label 16000
      Origin EGP, localpref 100, valid, internal, best, group-best, import-candidate,
not-in-vrf
      Received Path ID 0, Local Path ID 1, version 350584
      Extended community: RT:2:1
      Originator: 4.4.4.4, Cluster list: 3.3.3.3
```

The following is the sample output from the **show bgp** <*IP address*> command displaying the graceful-shutdown community and the graceful-shut path attribute with BGP graceful maintenance feature activated:

```
RP/0/0/CPU0:R4#show bgp 5.5.5.5
...
    10.10.10.1 from 10.10.10.1 (192.168.0.5)
      Received Label 24000
      Origin incomplete, metric 0, localpref 100, valid, internal, best, group-best,
```

```
import-candidate
      Received Path ID 0, Local Path ID 1, version 4
      Community: graceful-shutdown
      Originator: 192.168.0.5, Cluster list: 192.168.0.1
...
```

The following is sample output from the **show bgp** command in XR EXEC mode:

```
RP/0/RP0/CPU0:router#show bgp
 BGP router identifier 172.20.1.1, local AS number 1820
 BGP generic scan interval 60 secs
 BGP table state: Active
 Table ID: 0xe0000000
 BGP main routing table version 3
 Dampening enabled
 BGP scan interval 60 secs

 Status codes: s suppressed, d damped, h history, * valid, > best
               i - internal, S stale
 Origin codes: i - IGP, e - EGP, ? - incomplete

 Network            Next Hop            Metric LocPrf Weight Path
 * i10.3.0.0/16     172.20.22.1              0    100      0 1800 1239 ?
 *>i                172.20.16.1              0    100      0 1800 1239 ?
 * i10.6.0.0/16     172.20.22.1              0    100      0 1800 690 568 ?
 *>i                172.20.16.1              0    100      0 1800 690 568 ?
 * i10.7.0.0/16     172.20.22.1              0    100      0 1800 701 35 ?
 *>i                172.20.16.1              0    100      0 1800 701 35 ?
 *                  192.168.40.24                         0 1878 704 701 35 ?
 * i10.8.0.0/16     172.20.22.1              0    100      0 1800 690 560 ?
 *>i                172.20.16.1              0    100      0 1800 690 560 ?
 *                  192.168.40.24                         0 1878 704 701 560 ?
 * i10.13.0.0/16    172.20.22.1              0    100      0 1800 690 200 ?
 *>i                172.20.16.1              0    100      0 1800 690 200 ?
 *                  192.168.40.24                         0 1878 704 701 200 ?
 * i10.15.0.0/16    172.20.22.1              0    100      0 1800 174 ?
 *>i                172.20.16.1              0    100      0 1800 174 ?
 * i10.16.0.0/16    172.20.22.1              0    100      0 1800 701 i
 *>i                172.20.16.1              0    100      0 1800 701 i
 *                  192.168.40.24                         0 1878 704 701 i

 Processed 8 prefixes, 8 paths
```

This table describes the significant fields shown in the display.

Table 4: show bgp Field Descriptions

Field	Description
BGP router identifier	BGP identifier for the local system.
local AS number	Autonomous system number for the local system. • Range for 2-byte Autonomous system numbers (ASNs) is 1 to 65535. • Range for 4-byte Autonomous system numbers (ASNs) in asplain format is 1 to 4294967295. • Range for 4-byte Autonomous system numbers (ASNs) in asdot format is 1.0 to 65535.65535.

Field	Description
BGP generic scan interval	Interval (in seconds) between scans of the BGP table by a generic scanner.
BGP table state	State of the BGP database.
Table ID	BGP database identifier.
BGP main routing table version	Last version of the BGP database that was installed into the main routing table.
Dampening enabled	Dampening is enabled for the routes in this BGP routing table.
BGP scan interval	Interval (in seconds) between BGP scans for the specified address family and subaddress family.
Status codes	Status of the table entry. The status is displayed as a three-character field at the beginning of each line in the table. The first character may be (in order of precedence): S—Path is stale, indicating that a graceful restart is in progress with the peer from which the route was learned. s—Path is more specific than a locally sourced aggregate route and has been suppressed. *—Path is valid. The second character may be (in order of precedence): >—Path is the best path to use for that network. d—Path is dampened. h—Path is a history entry, representing a route that is currently withdrawn, but that is being maintained to preserve dampening information. Such routes should never be marked as valid. The third character may be: i—Path was learned by an internal BGP (iBGP) session.
Origin codes	Origin of the path. The origin code is displayed at the end of each line in the table. It can be one of the following values: i—Path originated from an Interior Gateway Protocol (IGP) and was sourced by BGP using a network or **aggregate-address** command. e—Path originated from an Exterior Gateway Protocol (EGP). ?—Origin of the path is not clear. Usually, this is a route that is redistributed into BGP from an IGP.
Network	IP prefix and prefix length for a network.
Next Hop	IP address of the next system that is used when a packet is forwarded to the destination network. An entry of 0.0.0.0 indicates that the router has a non-BGP route to this network.

Field	Description
Metric	Value of the interautonomous system metric, otherwise known as the Multi Exit discriminator (MED) metric.
LocPrf	Local preference value. This is used to determine the preferred exit point from the local autonomous system. It is propagated throughout the local autonomous system.
Weight	Path weight. Weight is used in choosing the preferred path to a route. It is not advertised to any neighbor.
Path	Autonomous system path to the destination network. At the end of the path is the path origin code.

The following is sample output from the **show bgp** command with the network specified:

```
RP/0/RP0/CPU0:router# show bgp 11.0.0.0/24
BGP router table entry for 11.0.0.0/24
 Versions:
  Process          bRIB/RIB  SendTblVer
  Speaker                2           2
 Last Modified: Mar  3 16:12:07.147 for 2d21h
 Paths: (3 available, best #1)
   Advertised to update-groups (with more than one peer):
     0.1
   Advertised to peers (in unique update groups):
     10.4.101.1
   Received by speaker 0
   Local
     0.0.0.0 from 0.0.0.0 (10.4.0.1)
       Origin IGP, metric 0, localpref 100, weight 32768, valid, local, best
   Received by speaker 0
   2 3 4
     10.4.101.1 from 10.4.101.1 (10.4.101.1)
       Origin IGP, localpref 100, valid, external
   Received by speaker 0
   Local
     10.4.101.2 from 10.4.101.2 (10.4.101.2)
       Origin IGP, localpref 100, valid, internal
```

The following is sample output from the **show bgp** command with the *ip-address/prefix-length* **detail** options:

```
RP/0/RP0/CPU0:router# show bgp 51.0.0.0/24 detail

Sat Mar 14 00:37:14.109 PST PDT

BGP routing table entry for 51.0.0.0/24

Versions:

  Process           bRIB/RIB  SendTblVer

  Speaker                  3           3

    Flags: 0x3e1000, label_retention: not enabled

Last Modified: Mar 13 19:32:17.976 for 05:04:56
```

```
Paths: (1 available, best #1)

  Advertised to update-groups (with more than one peer):

    0.3 0.4 0.7 0.8

  Advertised to peers (in unique update groups):

    201.48.20.1

  Path #1: Received by speaker 0

  Flags: 0x1000003

  200 201

    213.0.0.6 from 213.0.0.6 (200.200.3.1)

      Origin IGP, localpref 100, valid, external, best
```

The following is sample output from the **show bgp** command with the additional paths received from:

```
BGP routing table entry for 51.0.1.0/24, Route Distinguisher: 2:1
Versions:
  Process           bRIB/RIB  SendTblVer
  Speaker                 63          63
    Flags: 0x040630f2
Last Modified: Nov 11 12:44:05.811 for 00:00:16
Paths: (3 available, best #2)
  Advertised to CE peers (in unique update groups):
    10.51.0.10
  Path #1: Received by speaker 0
  Flags: 0x3
  Not advertised to any peer
  111 111 111 111 111 111 111 111
    10.51.0.10 from 10.51.0.10 (11.11.11.11)
      Origin IGP, metric 0, localpref 100, valid, external
      Received Path ID 0, Local Path ID 0, version 0
      Extended community: RT:55:1
  Path #2: Received by speaker 0
  Flags: 0x5060007
  Advertised to CE peers (in unique update groups):
    10.51.0.10
  561 562 563 564 565
    13.0.6.50 from 13.0.6.50 (13.0.6.50)
      Received Label 16
      Origin IGP, localpref 100, valid, internal, best, group-best, import-candidate,
imported
      Received Path ID 0, Local Path ID 1, version 63
      Extended community: RT:55:1
  Path #3: Received by speaker 0
  Flags: 0x4060007
  Not advertised to any peer
  591 592 593 594 595
    13.0.9.50 from 13.0.9.50 (13.0.9.50)
      Received Label 16
      Origin IGP, localpref 100, valid, internal, backup, add-path, import-candidate,
imported
      Received Path ID 0, Local Path ID 4, version 63
      Extended community: RT:22:232 RT:55:1
```

This is sample output to explain 'import suspect' state and 'import-suspect' field in **show bgp** command output:

```
RP/0/RP0/CPU0:router#show bgp vpnv4 unicast rd 11:111 100.16.11.0/24
BGP routing table entry for 100.16.11.0/24, Route Distinguisher: 11:111
Versions:
  Process           bRIB/RIB  SendTblVer
  Speaker            1834195     1834195
Paths: (2 available, best #1)
  Advertised to update-groups (with more than one peer):
    0.1
  Path #1: Received by speaker 0
  11
    1:16.16.16.16 (metric 30) from 55.55.55.55 (16.16.16.16)
      Received Label 19602
     Origin incomplete, localpref 100, valid, internal, best, import-candidate, not-in-vrf,
 import suspect
      Extended community: RT:11:11
      Originator: 16.16.16.16, Cluster list: 55.55.55.55
  Path #2: Received by speaker 0
  11
    1:16.16.16.16 (metric 30) from 88.88.88.88 (16.16.16.16)
      Received Label 19602
      Origin incomplete, localpref 100, valid, internal, not-in-vrf, import suspect
      Extended community: RT:11:11
      Originator: 16.16.16.16, Cluster list: 88.88.88.88
```

The **show bgp** command output displays 'import suspect' when potential import oscillation has been detected for the prefix. Import of such a prefix is not affected. However, import of the prefix can be dampened in future if the oscillation continues. If the oscillation stops during the next import run, the prefix will no longer be marked 'import supect'.

This is sample output of **show bgp {ipv4 | vpnv4} unicast summary** when the **update wait-install** command was configured for an address family. The output displays the "RIBAckVer" field.

```
RP/0/RP0/CPU0:router#show bgp summary

BGP router identifier 10.1.1.2, local AS number 100
BGP generic scan interval 60 secs
BGP table state: Active
Table ID: 0xe0000000   RD version: 5
BGP main routing table version 5
BGP scan interval 60 secs

BGP is operating in STANDALONE mode.


Process  RcvTblVer     RIBVer/RIBAckVer    LabelVer  ImportVer SendTblVer StandbyVer
Speaker          5          5/5                  5          5          5          5

Neighbor        Spk    AS MsgRcvd MsgSent   TblVer  InQ OutQ  Up/Down  St/PfxRcd
10.1.1.1          0   500       0       0        0    0    0 00:00:00 Idle
```

This is sample output from **show bgp vpnv4 unicast rd prefix/length** command that displays Accept Own prefix information:

```
RP/0/RP0/CPU0:router#show bgp vpnv4 unicast rd 10.10.10.10:1 110.1.1.1/32 detail
BGP routing table entry for 110.1.1.1/32, Route Distinguisher: 10.10.10.10:1
Versions:
```

```
  Process           bRIB/RIB  SendTblVer
  Speaker            1412487     1412487
    Local Label: 137742 (no rewrite);
    Flags: 0x04043001+0x00000000;
Last Modified: Jul 19 14:42:43.690 for 00:56:34
Paths: (2 available, best #1)
  Advertised to peers (in unique update groups):
    45.1.1.1
  Path #1: Received by speaker 0
  Flags: 0xd040003, import: 0x1f
  Advertised to peers (in unique update groups):
    45.1.1.1
  101
    10.5.1.2 from 10.5.1.2 (10.5.1.2)
      Origin incomplete, localpref 100, valid, external, best, group-best, import-candidate

      Received Path ID 0, Local Path ID 1, version 1412487
      Extended community: RT:100:1
  Path #2: Received by speaker 0
  Flags: 0x324020005, import: 0x01
  Not advertised to any peer
  101
    15.1.1.1 from 55.1.1.1 (15.1.1.1)
      Received Label 137742
      Origin incomplete, localpref 100, valid, internal, import-candidate, not-in-vrf,
accept-own-self
      Received Path ID 0, Local Path ID 0, version 0
      Community: accept-own
      Extended community: RT:100:1 RT:1000:1
      Originator: 15.1.1.1, Cluster list: 55.1.1.1, 75.1.1.1, 45.1.1.1
```

This is sample output from **show bgp vrf** *vrf-name* **ipv4unicast** *prefix/length* command that displays Accept Own prefix information on a customer (originating) VRF:

```
RP/0/RP0/CPU0:router#show bgp vrf customer1 ipv4 uni 110.1.1.1/32
BGP routing table entry for 110.1.1.1/32, Route Distinguisher: 10.10.10.10:1
Versions:
  Process           bRIB/RIB  SendTblVer
  Speaker            1412487     1412487
    Local Label: 137742
Last Modified: Jul 19 14:42:43.690 for 01:01:22
Paths: (2 available, best #1)
  Advertised to PE peers (in unique update groups):
    45.1.1.1
  Path #1: Received by speaker 0
  Advertised to PE peers (in unique update groups):
    45.1.1.1
  101
    10.5.1.2 from 10.5.1.2 (10.5.1.2)
      Origin incomplete, localpref 100, valid, external, best, group-best, import-candidate

      Received Path ID 0, Local Path ID 1, version 1412487
      Extended community: RT:100:1
  Path #2: Received by speaker 0
  Not advertised to any peer
  101
    15.1.1.1 from 55.1.1.1 (15.1.1.1)
      Received Label 137742
      Origin incomplete, localpref 100, valid, internal, import-candidate, not-in-vrf,
accept-own-self
      Received Path ID 0, Local Path ID 0, version 0
      Community: accept-own
```

```
      Extended community: RT:100:1 RT:1000:1
      Originator: 15.1.1.1, Cluster list: 55.1.1.1, 75.1.1.1, 45.1.1.1
```

This is sample output from **show bgp vrf** *vrf-name* **ipv4unicast** *prefix/length* command that displays Accept Own prefix information on a service VRF:

```
RP/0/RP0/CPU0:router#show bgp vrf service1 ipv4 uni 110.1.1.1/32
BGP routing table entry for 110.1.1.1/32, Route Distinguisher: 11.11.11.11:1
Versions:
  Process           bRIB/RIB  SendTblVer
  Speaker            1412497     1412497
Last Modified: Jul 19 14:43:08.690 for 01:39:22
Paths: (1 available, best #1)
  Advertised to CE peers (in unique update groups):
    10.8.1.2
  Path #1: Received by speaker 0
  Advertised to CE peers (in unique update groups):
    10.8.1.2
  101
    10.5.1.2 from 55.1.1.1 (15.1.1.1)
     Origin incomplete, localpref 100, valid, internal, best, group-best, import-candidate,
 imported, accept-own
      Received Path ID 0, Local Path ID 1, version 1412497
      Community: accept-own
      Extended community: RT:100:1 RT:1000:1
      Originator: 15.1.1.1, Cluster list: 55.1.1.1, 75.1.1.1, 45.1.1.1
```

This table describes the significant fields shown in the display:

Field	Description
accept-own-self	The Accept Own path in the customer VRF contains the "accept-own-self" keyword/flag.
accept-own	The Accept Own path contains the "accept-own" keyword/flag.
Community:accept-own	List of communities attached to the path: accept-own.
Extended community	List of extended communities attached to the path.
Cluster list	Router ID or cluster ID of all route reflectors through which the route has passed.

The output of **show bgp {vpnv4 | vpnv6} unicast rd** command may display the optional BGP attribute `not-in-vrf`. If a path in a VPNvX net is marked as not-in-vrf, it may be due to any of the following conditions:

- The RD of the VPNvX net is not the same as any of the RDs configured for VRFs on the router.
- The RD of the VPNvX net is the same as the RD configured for a specific VRF on the router, but the path is not imported to the specified VRF. For example, the route-targets attached to the path do not match any of the **import route-target** [*as-number:nn* | *ip-address:nn*] configured for VRF, *vrf_1*.

If the not-in-vrf net is set, it indicates that the path does not belong to the VRF.

This is sample output from the **show bgp ipv4 unicast** command showing the status of the permanent network:

```
RP/0/RP0/CPU0:router# show bgp ipv4 unicast 1.0.0.0/24
BGP routing table entry for 1.0.0.0/24
Versions:
  Process           bRIB/RIB  SendTblVer
  Speaker              90113       90113
Last Modified: Sep  6 04:46:03.650 for 00:14:19
Permanent Network
Paths: (2 available, best #2)
  Advertised to peers (in unique update groups):
    2.2.2.2
  Path #1: Received by speaker 0
  Advertised to peers (in unique update groups):
    3.3.3.3
  Local
    0.0.0.0 from 0.0.0.0 (1.1.1.1)
      Origin incomplete, metric 0, localpref 100, local, permanent-path
      Received Path ID 0, Local Path ID 4, version 90113
      Origin-AS validity: not-found
Path #2: Received by speaker 0
  Advertised to peers (in unique update groups):
    2.2.2.2
  7813 7814
    11.11.22.22 from 11.11.22.22 (192.1.1.1)
      Origin EGP, localpref 100, valid, external, best, group-best, import-candidate
      Received Path ID 0, Local Path ID 1, version 4
      Origin-AS validity: not-found
```

Examples

This is a sample output of **show bgp** command that displays large-communities in the list of attributes:

```
RP/0/RP0/CPU0:router#show bgp 3.3.3.3/32
Thu Mar 23 14:36:15.301 PDT
BGP routing table entry for 3.3.3.3/32
Versions:
  Process           bRIB/RIB  SendTblVer
  Speaker                 42          42
Last Modified: Mar 22 20:04:46.000 for 18:31:30
Paths: (1 available, best #1)
  Advertised to peers (in unique update groups):
    11.11.11.5
  Path #1: Received by speaker 0
  Advertised to peers (in unique update groups):
    11.11.11.5
  Local
    10.10.10.3 from 10.10.10.3 (3.3.3.3)
      Origin incomplete, metric 0, localpref 94, valid, internal, best, group-best
      Received Path ID 0, Local Path ID 0, version 42
      Community: 258:259 260:261 262:263 264:265
     Large Community: 1:2:3 5:6:7 4123456789:4123456780:4123456788
```

This example shows how to verify the RT-constraint default-prefix state for neighbor.

```
Router# show bgp ipv4 rt-filter neighbor 11.11.11.3 | include "RT-Constraint Default"
RT-Constraint Default Prefix: Suppressed-By-Config
```

show bgp bmp

To display Border Gateway Protocol (BGP) Monitoring Protocol (BMP) information, use the **show bgp bmp** command in XR EXEC mode.

show bgp bmp {**server** *server-id* **[detail]** | **summary**}

Syntax Description

server *server-id*	Displays information about BMP server as specified by the *server-id* variable.
detail	(Optional) Displays detailed BMP server information.
summary	Displays summary information about all the configured BMP servers.

Command Default

No default behavior or values

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

Task ID

Task ID	Operations
bgp	read

Examples

The following example shows sample output from the **show bgp bmp** command when the **summary** keyword is used:

```
RP/0/RP0/CPU0:router# show bgp bmp summary
ID   Host                Port     State   Time        NBRs
 1   10.0.101.1          16666    ESTAB   00:29:52    1
 2   10.0.101.2          16667    ESTAB   00:29:52    0
 3   fed0::1001          26666    ESTAB   00:29:52    0
 4   fed0::1002          26667    ESTAB   00:29:52    0
 5   10.0.101.1          16666    ESTAB   00:21:49    0
 6   10.0.101.1          16666    ESTAB   00:29:52    0
 7   fed0::1001          26666    ESTAB   00:29:52    0
 8   fed0::1001          26666    ESTAB   00:29:52    0
```

The following example shows sample output from the **show bgp bmp** command when the **server** keyword, with server ID as 4, is used:

```
RP/0/RP0/CPU0:router# show bgp bmp server 4
BMP server 4
Host 10.0.101.1 Port 16666
Connected for 00:25:07
```

```
Precedence:  internet
BGP neighbors: 1
VRF: - (0x60000000)
Update Source: 9.9.9.9 (Lo9)
Update Source Vrf ID: 0x60000000

Message Stats:
Total messages sent: 60
         INITIATION: 1
        TERMINATION: 0
       STATS-REPORT: 0
  PER-PEER messages: 59

  Neighbor 20.0.101.11
Messages pending: 0
Messages sent: 59
      PEER-UP: 1
    PEER-DOWN: 0
    ROUTE-MON: 58
```

show bgp update out

To display address-family level update generation information, use the **show bgp update out** command in XR EXEC mode.

show bgp [**vrf** *vrf-name*] [*afi safi*] **update out** [**brief** | **detail**]

Syntax Description

vrf *vrf-name*	(Optional) Displays non-default VRF.
aft	(Optional) Displays address-family identifier.
saft	(Optional) Displays subsequent address family identifier.
brief	(Optional) Displays brief information on process level update generation.
detail	(Optional) Displays detailed information on process level update generation.

Command Default

None

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 7.9.1	The command displays the summary of the neighbor address-family update-group, sub-group, or refresh sub-group information.
Release 6.0	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

Task ID

Task ID	Operation
bgp	read

This example displays sample output from the **show bgp update out** command:

```
RP/0/RP0/CPU0:router#show bgp update out
Address-family "IPv4 Unicast"
  Update generation status: Normal
  Update OutQ:              0 bytes (0 messages)
  AF update limit:  268435456 bytes (configured 268435456 bytes)
  EBGP Sub-group update limit:   33554432 bytes (configured 33554432 bytes)
  IBGP Sub-group update limit:   33554432 bytes (configured 33554432 bytes)

  Main routing table version: 2
  RIB version: 2
  Minimum neighbor version: 2
  AF Flags: 0x00000000
  Update-groups: 1
  Sub-groups: 1 (0 throttled)
```

```
  Refresh sub-groups: 0 (0 throttled)
  Filter-groups: 1
  Neighbors: 3

  History:
    Update OutQ Hi:                   300 bytes (1 messages)
    Update OutQ Cumulative:           600 bytes (2 messages)
    Update OutQ Discarded:              0 bytes (0 messages)
    Update OutQ Cleared:                0 bytes (0 messages)
    Last discarded from OutQ:  --- (never)
    Last cleared   from OutQ:  --- (never)
    Update generation throttled 0 times, last event --- (never)
    Update generation recovered 0 times, last event --- (never)
    Update generation mem alloc failed 0 times, last event --- (never)

  VRF "default", Address-family "IPv4 Unicast"
    RD flags: 0x00000001
    RD Version: 2
    Table flags: 0x00000021
    RIB version: 2
    Update-groups: 1
    Sub-groups: 1 (0 throttled)
    Refresh sub-groups: 0 (0 throttled)
    Filter-groups: 1
    Neighbors: 3

RP/0/RSP0/CPU0:PE51_ASR-9010#
RP/0/RSP0/CPU0:PE51_ASR-9010#
RP/0/RSP0/CPU0:PE51_ASR-9010#show bgp update out filter-group
Thu Sep 13 01:43:48.183 DST
```

The command shows summary of the neighbor address-family update-group, sub-group, or refresh sub-group information. It is modified to show if the peer is configured as *Static* slow peer or *Dynamic* slow peer. The command is also modified to show summary for the slow peers only.

The **show bgp ipv6 unicast update out neighbor brief** command displays output for all peers.

The **show bgp ipv6 unicast update out neighbor slow-peers brief** commnad displays the output for the slow peers only.

```
Router#show bgp ipv6 unicast update out neighbor slow-peers brief

Fri Nov 18 04:53:32.903 UTC

VRF "default", Address-family "IPv6 Unicast"
  Main routing table version: 1572958
  RIB version: 1572958

Legend: (S) - Slow peer static configured
        (D) - Slow peer dynamic detected

  Neighbor        FG      SG      SG-R           UG      Status OutQ       OutQ-R     Version
     Ack/Ack-R

  2020:102::1     0.52    0.24    ---            0.18    Normal 0          0          1572958
     1572958 (S)
  2020:103::1     0.53    0.25    ---            0.19    Normal 0          0          1572958
     1572958 (S)
  2020:104::1     0.8     0.8     ---            0.20    Normal 0          0          1572958
     1572958 (S)
........
  2020:11c::1     0.43    0.14    0.14:1853      0.4     Normal 0          0          1572958
     1572958/0 (D)
```

```
  2020:129::1     0.3     0.3     0.3:1833       0.6     Normal 0          0          1572958
     1572958/0 (D)
  2020:149::1     0.3     0.3     0.3:1849       0.6     Normal 0          0          1572958
     1572678/0 (D)
  2020:156::1     0.43    0.14    0.14:1854      0.4     Normal 0          0          1572958
     1572958/0 (D)
---More----
```

The command displays all neighbors update information. The output is enhanced to include slow peer details. The command has also been modified to show neighbors update information of only slow peer.

```
Router#show bgp update out neighbor slow-peers detail
Wed Jun 1 13:34:23.605 IST

VRF "default", Address-family "IPv4 Unicast"
Main routing table version: 47521
RIB version: 47521

Neighbor 192.168.0.4
Filter-group 0.1, Refresh filter-group 0.1
Sub-group 0.1, Refresh sub-group 0.1:1
Update-group 0.2

Update OutQ: 30000 bytes (50 messages)
Refresh update OutQ: 10800 bytes (18 messages)
Filter-group pending: 68 messages

Neighbor flags: 0x40310060+0x00002026
Reset 0x00003fef, Sync 0x00000000, Ver catchup 0x00000000
GSHUT 0x00000000
Neighbor AF flags: 0x00000204+0x00020030+0x00280000+0x00000018
Capability 0x00000801, GR 0x00000000

Version: 47521
Ack version: 25021 (Synced 0)
Outstanding version count: 8
Pending target version: 0 (next resume: 0)
EOR outstanding ? [No]
Refresh version: 31521
Refresh Ack version: 0
Refresh target version: 31521 (requested 31521)
Refresh pending target version: 0 (next resume: 0)
Refresh State: SLOW-RTC
Update Message Pointer: 0x7f17d4ae6988
Last Message Enqueued: 4 secs
Slow Detection State: Dynamic Detected Slow Peer
----More-----
```

The **show bgp update out neighbor slow-peers detail** command displays the output of a specific neighbor address-family for the slow peers only.

```
Router#show bgp vpnv4 unicast update out neighbor 20.1.114.1 detail

Tue Nov 15 18:48:24.863 UTC

VRF "default", Address-family "VPNv6 Unicast"
  Main routing table version: 51001
  RIB version: 51001

  Neighbor 20.1.114.1
    Filter-group 0.50,   Refresh filter-group 0.50
```

```
    Sub-group 0.5,   Refresh sub-group 0.5:538
    Update-group 0.2

    Update OutQ:              147400 bytes (67 messages)
    Refresh update OutQ:      341000 bytes (155 messages)
    Filter-group pending:                  222 messages

    Neighbor flags: 0x40218060+0x00000026
      Reset 0x00003fef,   Sync 0x00000000,   Ver catchup 0x00000000
      GSHUT 0x00000000,   NbrInfo 0x00000044
    Neighbor AF flags: 0x00110224+0x10060020+0x00000000+0x00000048
      Capability 0x00000001,   GR 0x00000000
      DR 0x00000000,   DR In use 0x00000000
      LLGR 0x00000000

    Version: 51001
    Ack version: 25750 (Synced 0)
    Outstanding version count: 0
    Pending target version: 0 (next resume: 0)
    EOR outstanding ?  [No]
    Refresh version: 25750
    Refresh Ack version: 0
    Refresh target version: 25750 (requested 50251)
    Refresh pending target version: 0 (next resume: 0)
    Refresh State: SLOW
    Update Message Pointer: 0x7f17d4bc6918
    Last Message Enqueued: 20 secs
    Slow Detection State: Dynamic Detected Slow Peer
----More----
```

Where:

- *Refresh State* indicates the state of the refresh sub-group. The states can be any of the following:
  - *Not-In-Refresh*: Refresh sub-group is not present
  - *RR*: Refresh sub-group is processing refresh request update
  - *SLOW*: Refresh sub-group is processing slow peer update
  - *RTC*: Refresh sub-group is processing RTC incremental update
  - *SLOW-RTC*: Refresh sub-group is processing both slow peer and RTC incremental update
- *Last Message Enqueued* indicates the time since the last update message was enqueued to the neighbor address-family. If this time exceeds the neighbor address-family slow peer detection threshold time, then it meets one of the conditions for detection of slow peer.
- *Slow Detection State* indicates the operational type of slow peer. The type can be any of the following:
  - *Static Slow Peer*: Neighbor address-family is static slow peer
  - *Dynamic Detected Slow Peer*: Neighbor address-family is dynamic detected slow peer
  - *Not slow peer*: Neighbor address-family is not a slow peer

show bgp update in error process

To display process level update inbound error-handling information, use the **show bgp update in error process**command in XR EXEC mode.

show bgp update in error process [**brief** | **detail**]

Syntax Description

brief	(Optional) Displays brief information on process level update generation.
detail	(Optional) Displays detailed information on process level update generation.

Command Default

None

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

Task ID

Task ID	Operations
bgp	read

This example displays sample output from the **show bgp update in error process** command:

```
RP/0/RP0/CPU0:router#show bgp update in error process

Basic Update error-handling:
  EBGP: [Enabled]
  IBGP: [Enabled]
Extended Update error-handling:
  EBGP: [Disabled]
  IBGP: [Disabled]

Malformed Update messages: 0
Neighbors that received malformed Update messages: 0
Last malformed Update received: --- (never)
```

show bgp update out filter-group

To display update generation information at filter-group level, **show bgp update out filter-group** command in XR EXEC mode.

show bgp [**vrf** *vrf-name*] [*afi safi*] **update out filter-group** [*fg-process-id*] [**brief** | **detail**]

Syntax Description

vrf *vrf-name*	Specifies the non-default VRF.
afi safi	Specifies the address family and subsequent address family identifiers.
fg-process-id	Specifies the filter-group process ID in <x.y> format. Range is < 0-15>.<0-4294967295>.
brief	(Optional) Displays brief information on filter-group level update generation
detail	(Optional) Displays detailed information on filter-group level update generation.

Command Default

None

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

Task ID

Task ID	Operations
bgp	read

This example displays sample output from **show bgp update out filter-group** command:

show bgp update out process

To display process level update generation information, use the **show bgp update out process** command in XR EXEC mode.

show bgp update out process [**brief** | **detail**]

Syntax Description

brief	(Optional) Displays brief information on process level update generation.
detail	(Optional) Displays detailed information on process level update generation.

Command Default

None

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

Task ID

Task ID	Operation
bgp	read

This example displays sample output from the show bgp update out process brief command:

```
RP/0/RP0/CPU0:router#show bgp update out process
Wed Sep 12 08:26:04.308 DST

Update generation status: Normal
Update OutQ:           0 bytes (0 messages)
Update limit:  536870912 bytes (configured 536870912 bytes)

Update generation logging:  [Disabled]

  Address-family Status    Limit      OutQ       UG   SG(Thr)    SG-R(Thr) Nbrs

  IPv4 Unicast   Normal    268435456  0          1    1(0)       0(0)       3
  L2VPN VPLS     Normal    268435456  0          1    1(0)       0(0)       3

History:
  Update OutQ Hi:                  300 bytes (1 messages)
  Update OutQ Cumulative:         1200 bytes (4 messages)
  Update OutQ Discarded:             0 bytes (0 messages)
  Update OutQ Cleared:               0 bytes (0 messages)
  Last discarded from OutQ:  --- (never)
  Last cleared   from OutQ:  --- (never)
  Update generation throttled 0 times, last event --- (never)
  Update generation recovered 0 times, last event --- (never)
```

```
Update generation mem alloc failed 0 times, last event --- (never)
```

show bgp update out sub-group

To display sub-group update generation information, use the **show bgp update out sub-group** command in XR EXEC mode.

show bgp [**vrf** *vrf-name*] [*afi safi*] **update out** [**update-group** *ug-index*] **sub-group** [*sg-index*] [**brief** | **detail**]

Syntax Description

vrf *vrf-name*	(Optional) Displays non-default VRF.
aft	(Optional) Displays address-family identifier.
saft	(Optional) Displays subsequent address family identifier.
brief	(Optional) Displays brief information on process level update generation.
detail	(Optional) Displays detailed information on process level update generation.
ug-index	(Optional) Displays the update-group process ID in <x.y> format.
sg-index	(Optional) displays the sub-group process ID in <x.y> format.

Command Default

None

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

Task ID

Task ID	Operation
bgp	read

This example displays sample output from the **show bgp update out sub-group** command:

```
RP/0/RP0/CPU0:router#show bgp update out sub-group

VRF "default", Address-family "IPv4 Unicast"
  Main routing table version: 2
  RIB version: 2

  SG             UG      Status    Limit      OutQ       SG-R Nbrs Version    ()

  0.2            0.2     Normal    33554432   0          0    3    2          ()
RP/0/RSP0/CPU0:PE51_ASR-9010#
```

This table describes the significant fields shown in the display:

show bgp update out update-group

To display update-group update generation information, use the **show bgp update out update-group** command in XR EXEC mode.

show bgp [**vrf** *vrf-name*] [*afi safi*] **update out update-group** [*ug-index*] [**brief** | **detail**]

Syntax Description

vrf *vrf-name*	(Optional) Displays non-default VRF.
aft	(Optional) Displays address-family identifier.
saft	(Optional) Displays subsequent address family identifier.
brief	(Optional) Displays brief information on process level update generation.
detail	(Optional) Displays detailed information on process level update generation.
ug-index	(Optional) Displays the update-group process ID in <x.y> format.

Command Default

None

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

Task ID

Task ID	Operation
bgp	read

This example shows the significant fields on display form the **show bgp update out update-group** command:

```
RP/0/RP0/CPU0:router#show bgp update out sub-group

VRF "default", Address-family "IPv4 Unicast"
  Main routing table version: 2
  RIB version: 2

  SG             UG      Status    Limit      OutQ       SG-R Nbrs Version    ()

  0.2            0.2     Normal    33554432   0          0    3    2          ()
RP/0/RSP0/CPU0:PE51_ASR-9010#show bgp update ou update-group
Wed Sep 12 08:37:24.756 DST

VRF "default", Address-family "IPv4 Unicast"
```

```
UG       OutQ       SG(Thr)     SG-R(Thr)   FG    Nbrs

0.2      0          1(0)        0(0)        1     3
```

show bgp vrf update in error

To display VRF level update inbound error-handling information, use the **show bgp vrf update in error** command in XR EXEC mode.

show bgp [**vrf** *vrf-name*] **update in error** [**brief** | **detail**]

vrf *vrf-name*	(Optional) Displays non-default VRF.
brief	(Optional) Displays brief information.
detail	(Optional) Displays detailed information.

Command Default

None

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

Task ID

Task ID	Operations
bgp	read

This example displays sample output from s**how bgp vrf vrf1 update in error** command:

```
RP/0/RP0/CPU0:router#show bgp update in error

VRF "default"
  Malformed Update messages: 0
  Neighbors that received malformed Update messages: 0
  Last malformed update received: --- (never)
```

show bgp advertised

To display advertisements for neighbors or a single neighbor, use the **show bgp advertised**command in XR EXEC mode.

show bgp [**ipv4** { **all** | **labeled-unicast** | **mdt** | **multicast** | **tunnel** | **unicast** }] **advertised** [**neighbor** *ip-address*] [**standby**] [**summary**]

show bgp [**ipv6** { **all** | **labeled-unicast** | **multicast** | **unicast**}] **advertised** [**neighbor** *ip-address*] [**standby**] [**summary**]

show bgp [**all** { **all** | **labeled-unicast** | **multicast** | **tunnel** | **unicast** }] **advertised** [**neighbor** *ip-address*] [**standby**] [**summary**]

show bgp [**vpnv4 unicast** [**rd** *rd-address*]] **advertised** [**neighbor** *ip-address*] [**standby**] [**summary**]

show bgp [**vpnv6 unicast** [**rd** *rd-address*]] **advertised** [**neighbor** *ip-address*] [**standby**] [**summary**]

show bgp [**vrf** {*vrf-name* | **all**} [**ipv4** | {**labeled-unicast** | **unicast**} | **ipv6 unicast**]] **advertised** [**neighbor** *ip-address*] [**standby**] [**summary**]

Syntax Description

ipv4	(Optional) Specifies IP Version 4 address prefixes.
unicast	(Optional) Specifies unicast address prefixes.
multicast	(Optional) Specifies multicast address prefixes.
labeled-unicast	(Optional) Specifies labeled unicast address prefixes.
all	(Optional) For address family, specifies prefixes for all address families.
tunnel	(Optional) Specifies tunnel address prefixes.
ipv6	(Optional) Specifies IP Version 6 address prefixes.
vpnv4 unicast	(Optional) Specifies VPNv4 unicast address families.
rd *rd-address*	(Optional) Displays routes with a specific route distinguisher.
vrf	(Optional) Specifies VPN routing and forwarding (VRF) instance.
vrf-name	(Optional) Name of a VRF.
all	(Optional) For VRF, specifies all VRFs.
ipv4 { **unicast** | **labeled-unicast** }	(Optional) For VRF, specifies IPv4 unicast or labeled-unicast address families.
ipv6 unicast	(Optional) For VRF, specifies IPv6 unicast address families.

neighbor	(Optional) Previews advertisements for a single neighbor. If the **neighbor** keyword is omitted, then the advertisements for all neighbors are displayed.
ip-address	(Optional) IP address of the neighbor.
summary	(Optional) Displays a summary of advertisements.

Command Default

If no address family or subaddress family is specified, the default address family and subaddress family specified using the **set default-afi** and **set default-safi** commands are used.

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

Note The **set default-afi** command is used to specify the default address family for the session, and the **set default-safi** command is used to specify the default subaddress family for the session. See the for detailed information and syntax for the **set default-afi** and **set default-safi** commands. If you do not specify a default address family, the default address family is IPv4. If you do not specify a default subaddress family, the default subaddress family is unicast.

BGP contains a separate routing table for each address family and subaddress family combination that is configured. The address family and subaddress family options specify the routing table to be examined. If the **all** keyword is specified for the address family or subaddress family, each matching routing table is examined in turn.

Use the **show bgp advertised** command to display the routes that have been advertised to peers or a specific peer. To preview advertisements that would be sent to a peer under a particular policy, even if the corresponding update messages have not been generated yet, use the **show bgp policy** command.

Note When you issue the **show bgp advertised** command, a route is not displayed in the output unless an advertisement for that route has already been sent (and not withdrawn). If an advertisement for the route has not yet been sent, the route is not displayed.

Use the **summary** keyword to display a summary of the advertised routes. If you do not specify the **summary** keyword, the software displays detailed information about the advertised routes.

Note The **show bgp advertised** command does not display the application of any outbound policy in the route details it displays. Consequently, this command provides only an indication of whether a particular route has been advertised, rather than details of which attributes were advertised. Use the **show bgp policy sent-advertisements** command to display the attributes that are advertised.

Warning The **show bgp advertised** command for any AFI/SAFI can increase CPU utilization under scale conditions. At high prefix scale, high neighbor scale, or both, the command can respond slowly and degrade concurrent BGP show operations.

Task ID

Task ID	Operations
bgp	read

Examples

The following is sample output from the **show bgp advertised** command in XR EXEC mode:

```
RP/0/RP0/CPU0:router# show bgp advertised neighbor 10.0.101.4 summary

Network     Next Hop            From          AS Path
1.1.1.0/24  10.0.101.1          10.0.101.1    2 3 222 333 444 555 i
1.1.2.0/24  10.0.101.1          10.0.101.1    3 4 5 6 7 i
1.1.3.0/24  10.0.101.1          10.0.101.1    77 88 33 44 55 99 99 99 i
1.1.4.0/24  10.0.101.1          10.0.101.1    2 5 6 7 8 i
1.1.7.0/24  10.0.101.1          10.0.101.1    3 5 i
1.1.8.0/24  10.0.101.1          10.0.101.1    77 88 99 99 99 i
```

This table describes the significant fields shown in the display.

Table 5: show bgp advertised neighbor summary Field Descriptions

Field	Description
Network	IP prefix and prefix length for a network.
Next Hop	IP address of the next system that is used when a packet is forwarded to the destination network. An entry of 0.0.0.0 indicates that the router has a non-BGP route to this network.
From	IP address of the peer that advertised this route.
AS Path	AS path of the peer that advertised this route.
Local	Indicates the route originated on the local system.
Local Aggregate	Indicates the route is an aggregate created on the local system.

Field	Description
Advertised to	Indicates the peer to which this entry was advertised. This field is used in the output when displaying a summary of the advertisements to all neighbors.

The following is sample output from the **show bgp advertised** command for detailed advertisement information:

```
RP/0/RP0/CPU0:router# show bgp advertised neighbor 172.72.77.1

172.16.0.0/24 is advertised to 172.72.77.1
  Path info:
    neighbor: Local            neighbor router id: 172.74.84.1
    valid  redistributed  best
  Attributes after inbound policy was applied:
 next hop: 0.0.0.0
    MET ORG AS
    origin: incomplete  metric: 0
    aspath:
10.52.0.0/16 is advertised to 172.72.77.1
  Path info:
    neighbor: Local Aggregate  neighbor router id: 172.74.84.1
    valid  aggregated  best
  Attributes after inbound policy was applied:
 next hop: 0.0.0.0
    ORG AGG ATOM
    origin: IGP  aggregator: 172.74.84.1 (1)
    aspath:
```

This table describes the significant fields shown in the display.

Table 6: show bgp advertised neighbor Field Descriptions

Field	Description
is advertised to	IP address of the peer to which this route has been advertised. If the route has been advertised to multiple peers, the information is shown separately for each peer.
neighbor	IP address of the peer that advertised this route, or one of the following: Local—Route originated on the local system. Local Aggregate—Route is an aggregate created on the local system.
neighbor router id	BGP identifier for the peer, or the local system if the route originated on the local system.
Not advertised to any peer	Indicates the no-advertise well-known community is associated with this route. Routes with this community are not advertised to any BGP peers.
Not advertised to any EBGP peer	Indicates the no-export well-known community is associated with this route. Routes with this community are not advertised to external BGP peers, even if those external peers are part of the same confederation as the local router.
Not advertised outside the local AS	Indicates the local-AS well-known community is associated with this route. Routes with this community value are not advertised outside the local autonomous system or confederation boundary.

Field	Description
(Received from a RR-client)	Path was received from a route reflector client.
(received-only)	This path is not used for routing purposes. It is used to support soft reconfiguration, and records the path attributes before inbound policy was applied to a path received from a peer. A path marked "received-only" indicates that either the path was dropped by inbound policy, or the path information was modified by inbound policy and a separate copy of the modified path is used for routing.
(received & used)	Indicates that the path is used both for soft reconfiguration and routing purposes. A path marked "received and used," implies the path information was not modified by inbound policy.
valid	Path is valid.
redistributed	Path is locally sourced through redistribution.
aggregated	Path is locally sourced through aggregation.
local	Path is locally sourced through the **network** command.
confed	Path was received from a confederation peer.
best	Path is selected as best.
multipath	Path is one of multiple paths selected for load-sharing purposes.
dampinfo	Indicates dampening information: Penalty—Current penalty for this path. Flapped—Number of times the route has flapped. In—Time (hours:minutes:seconds) since the router noticed the first flap. Reuse in—Time (hours:minutes:seconds) after which the path is made available. This field is displayed only if the path is currently suppressed.

Field	Description
Attributes after inbound policy was applied	Displays attributes associated with the received route, after any inbound policy has been applied. AGG—Aggregator attribute is present. AS—AS path attribute is present. ATOM—Atomic aggregate attribute is present. COMM—Communities attribute is present. EXTCOMM—Extended communities attribute is present. LOCAL—Local preference attribute is present. MET—Multi Exit Discriminator (MED) attribute is present. next hop—IP address of the next system used when a packet is forwarded to the destination network. An entry of 0.0.0.0 indicates that the router has a non-BGP route to this network. ORG—Origin attribute is present.
origin	Origin of the path: IGP—Path originated from an Interior Gateway Protocol (IGP) and was sourced by BGP using a **network** or **aggregate-address** command. EGP—Path originated from an Exterior Gateway Protocol. incomplete—Origin of the path is not clear. For example, a route that is redistributed into BGP from an IGP.
neighbor as	First autonomous system (AS) number in the AS path.
aggregator	Indicates that the path was received with the aggregator attribute. The autonomous system number and router-id of the system that performed the aggregation are shown.
metric	Value of the interautonomous system metric, otherwise known as the MED metric.
localpref	Local preference value. This is used to determine the preferred exit point from the local autonomous system. It is propagated throughout the local autonomous system
aspath	AS path associated with the route.

Field	Description
community	Community attributes associated with the path. Community values are displayed in AA:NN format, except for the following well-known communities: Local-AS—Community with value 4294967043 or hex 0xFFFFFF03. Routes with this community value are not advertised outside the local autonomous system or confederation boundary. no-advertise—Community with value 4294967042 or hex 0xFFFFFF02. Routes with this community value are not advertised to any BGP peers. no-export—Community with value 4294967041 or hex 0xFFFFFF01. Routes with this community are not advertised to external BGP peers, even if those peers are in the same confederation with the local router.
Extended community	Extended community attributes associated with the path. For known extended community types, the following codes may be displayed: RT—Route target community SoO—Site of Origin community LB—Link Bandwidth community
Originator	Router ID of the originating router when route reflection is used.
Cluster lists	Router ID or cluster ID of all route reflectors through which the route has passed.

show bgp af-group

To display information about Border Gateway Protocol (BGP) configuration for address family groups, use the **show bgp af-group** command in XR EXEC mode.

show bgp af-group *group-name* {**configuration** [**defaults**] [**nvgen**] | **inheritance** | **users**}

Syntax Description

group-name	Name of the address family group to display.
configuration	(Optional) Displays the effective configuration for the af-group, including any settings that have been inherited from af-groups used by this af-group.
defaults	(Optional) Displays all configuration settings, including any default settings.
nvgen	(Optional) Displays output in the format of **show running-config** output. If the **defaults** keyword is also specified, the output is not suitable for cutting and pasting into a configuration session.
inheritance	Displays the af-groups from which this af-group inherits configuration settings.
users	Displays the neighbors, neighbor groups, and af-groups that inherit configuration from this af-group.

Command Default

No default behavior or value

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

Use the **show bgp af-group** command with the *group-name* **configuration** argument and keyword to display the effective configuration of an af-group, taking into account any configuration that may be inherited from other af-groups through the **use af-group** command. The source of each command is shown.

If the **defaults** keyword is specified, all configuration for the af-group, including default values, is shown. Default configuration is identified in the show output. Use the **nvgen** keyword to display configuration formatted in the style of the **show running-config** command. This output is suitable for cutting and pasting into configuration sessions.

Use the **show bgp af-group** command with the *group-name* **inheritance** argument and keyword to display the address family groups from which the specified af-group inherits configuration.

Use the **show bgp af-group** command with the *group-name* **users** argument and keyword to display the neighbors, neighbor groups, and af-groups that inherit configuration from the specified af-group.

Task ID

Task ID	Operations
bgp	read

Examples

The following af-group configuration is used in the examples:

```
af-group group3 address-family ipv4 unicast
remove-private-AS
soft-reconfiguration inbound
!
af-group group1 address-family ipv4 unicast
use af-group group2
maximum-prefix 2500 75 warning-only
default-originate
soft-reconfiguration inbound disable
!
af-group group2 address-family ipv4 unicast
use af-group group3
send-community-ebgp
send-extended-community-ebgp
capability orf prefix both
```

The following is sample output from the **show bgp af-group** command with the **configuration** keyword in XR EXEC mode. The source of each command is shown in the right column. For example, **default-originate** is configured directly on **af-group group1** , and the **remove-private-AS** command is inherited from af-group group2, which in turn inherits it from af-group group3.

```
RP/0/RP0/CPU0:router# show bgp af-group group1 configuration

af-group group1 address-family ipv4 unicast
  capability orf prefix both                [a:group2]
  default-originate                         []
  maximum-prefix 2500 75 warning-only       []
  remove-private-AS                         [a:group2 a:group3]
  send-community                            [a:group2]
  send-extended-community                   [a:group2
```

The following is sample output from the **show bgp af-group** command with the **users** keyword:

```
RP/0/RP0/CPU0:router# show bgp af-group group2 users

IPv4 Unicast: a:group1
```

The following is sample output from the **show bgp af-group** command with the **inheritance** keyword. This example shows that the specified af-group group1 directly uses the group2 af-group, which in turn uses the group3 af-group:

```
RP/0/RSP0RP0/CPU0:router# show bgp af-group group1 inheritance

IPv4 Unicast: a:group2 a:group3
```

Table 7: show bgp af-group Field Descriptions, on page 305 describes the significant fields shown in the display.

This table describes the significant fields shown in the display.

Table 7: show bgp af-group Field Descriptions

Field	Description
[ ]	Configures the command directly on the specified address family group.
a:	Indicates the name that follows is an address family group.
n:	Indicates the name that follows is a neighbor group.
[dflt]	Indicates the setting is not explicitly configured or inherited, and the default value for the setting is used. This field may be shown when the **defaults** keyword is specified.
<not set>	Indicates that the configuration is disabled by default. This field may be shown when the **defaults** keyword is specified.

show bgp attribute-key

To display all existing attribute keys, use the **show bgp attribute-key**command in XR EXEC mode.

show bgp {**ipv4** | **ipv6** | **all** | **vpnv4 unicast** | }
{**vrf** | **vpnv6 unicast**} **attribute-key** [**standby**]

Syntax Description

ipv4	(Optional) Specifies IP Version 4 address prefixes.
unicast	(Optional) Specifies unicast address prefixes.
multicast	(Optional) Specifies multicast address prefixes.
labeled-unicast	(Optional) Specifies labeled unicast address prefixes.
all	(Optional) For address family, specifies prefixes for all address families.
tunnel	(Optional) Specifies tunnel address prefixes.
all	(Optional) For subaddress family, specifies prefixes for all subaddress families.
ipv6	(Optional) Specifies IP Version 6 address prefixes.
vpnv4-unicast	(Optional) Specifies VPNv4 unicast address families.
vrf	(Optional) Specifies VPN routing and forwarding (VRF) instance.
vrf-name	(Optional) Name of a VRF.
all	(Optional) For VRF, specifies all VRFs.
ipv4 { **unicast** | **labeled-unicast** }	(Optional) For VRF, specifies IPv4 unicast or labeled-unicast address families.

Command Default

If no address family or subaddress family is specified, the default address family and subaddress family specified using the **set default-afi** and **set default-safi** commands are used.

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines



Note The **set default-afi** command is used to specify the default address family for the session, and the **set default-safi** command is used to specify the default subaddress family for the session. See the for detailed information and syntax for the **set default-afi** and **set default-safi** commands. If you do not specify a default address family, the default address family is IPv4. If you do not specify a default subaddress family, the default subaddress family is unicast.

Task ID

Task ID	Operations
bgp	read

Examples

The following is sample output from the **show bgp attribute-key**command in XR EXEC mode:

```
RP/0/RP0/CPU0:router# show bgp all all attribute-key

Address Family: IPv4 Unicast
============================

BGP router identifier 10.0.0.1, local AS number 1
BGP generic scan interval 60 secs
BGP main routing table version 109
BGP scan interval 60 secs
Status codes: s suppressed, d damped, h history, * valid, > best
              i - internal, S stale
Origin codes: i - IGP, e - EGP, ? - incomplete
   Network            Next Hop        AttrKey
*> 1.1.0.0/16         0.0.0.0         0x00000002
*> 10.0.0.0/16        0.0.0.0         0x00000002
*> 12.21.0.0/16       0.0.0.0         0x00000002
*> 194.3.192.1/32     10.0.101.1      0x00000009
*> 194.3.192.2/32     10.0.101.1      0x00000009
*> 194.3.192.3/32     10.0.101.1      0x00000009
*> 194.3.192.4/32     10.0.101.1      0x00000009
*> 194.3.192.5/32     10.0.101.1      0x00000009

Processed 8 prefixes, 8 paths

Address Family: IPv4 Multicast
==============================

BGP router identifier 10.0.0.1, local AS number 1
BGP generic scan interval 60 secs
BGP main routing table version 15
BGP scan interval 60 secs
Status codes: s suppressed, d damped, h history, * valid, > best
              i - internal, S stale
Origin codes: i - IGP, e - EGP, ? - incomplete
   Network            Next Hop        AttrKey
*> 194.3.193.2/32     10.0.101.1      0x00000009
*> 194.3.193.3/32     10.0.101.1      0x00000009

Processed 2 prefixes, 2 paths

Address Family: IPv6 Unicast
```

```
============================

BGP router identifier 10.0.0.1, local AS number 1
BGP generic scan interval 60 secs
BGP main routing table version 19
BGP scan interval 60 secs
Status codes: s suppressed, d damped, h history, * valid, > best
              i - internal, S stale
Origin codes: i - IGP, e - EGP, ? - incomplete
   Network            Next Hop        AttrKey
*> 2222::1111/128     2222::2         0x00000009
*> 2222::1112/128     2222::2         0x00000009

Processed 2 prefixes, 2 paths
```

This table describes the significant fields shown in the display.

Table 8: show bgp attribute-key Field Descriptions

Field	Description
BGP router identifier	BGP identifier for the local system.
local AS number	Autonomous system number for the local system.
BGP generic scan interval	Interval (in seconds) between scans of the BGP table by a generic scanner.
BGP main routing table version	Last version of the BGP database that was installed into the main routing table.
BGP scan interval	Interval (in seconds) between scans.
Status codes	Status of the table entry. The status is displayed as a three-character field at the beginning of each line in the table. The first character may be (in order of precedence): S—Path is stale, indicating that a graceful restart is in progress with the peer from which the route was learned. s—Path is more specific than a locally sourced aggregate route and has been suppressed. *—Path is valid. The second character may be (in order of precedence): >—Path is the best path to use for that network. d—Path is dampened. h—Path is a history entry, representing a route that is currently withdrawn, but that is being maintained to preserve dampening information. Such routes should never be marked as valid. The third character may be: i—Path was learned by an internal BGP (iBGP) session.

Field	Description
Origin codes	Origin of the path. The origin code is displayed at the end of each line in the table. It can be one of the following values: i—Path originated from an Interior Gateway Protocol (IGP) and was sourced by BGP using a **network** or **aggregate-address** command. e—Entry originated from an Exterior Gateway Protocol (EGP). ?—Origin of the path is not clear. Usually, this is a route that is redistributed into BGP from an IGP.
Network	IP prefix and prefix length for a network.
Next Hop	IP address of the next system that is used when a packet is forwarded to the destination network. An entry of 0.0.0.0 indicates that the router has a non-BGP route to this network.
AttrKey	Key associated with the route attribute.
Processed *n* prefixes, *n* paths	Number of prefixes and number of paths processed for the table.

show bgp cidr-only

To display routes with nonnatural network masks, also known as classless interdomain routing (CIDR) routes, use the **show bgp cidr-only** command in XR EXEC mode.

show bgp [**ipv4** | **vrf**] **cidr-only** [**standby**]

Syntax Description

ipv4	(Optional) Specifies the IP Version 4 address family.
unicast	(Optional) Specifies the unicast address family.
multicast	(Optional) Specifies the multicast address family.
labeled-unicast	(Optional) Specifies labeled unicast address prefixes.
all	(Optional) For subaddress family, specifies all subaddress families.
tunnel	(Optional) Specifies the tunnel address family.
vrf	(Optional) Specifies VPN routing and forwarding (VRF) instance.
vrf-name	(Optional) Name of a VRF.
all	(Optional) For VRF, specifies all VRFs.
ipv4 { **unicast** | **labeled-unicast** }	(Optional) For VRF, specifies IPv4 unicast or labeled-unicast address families.

Command Default

If no address family or subaddress family is specified, the default address family and subaddress family specified using the **set default-afi** and **set default-safi** commands are used. This command is applicable only for IPv4 prefixes. If the default address family is not IPv4, then the **ipv4** keyword must be used.

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

Note The **set default-afi** command is used to specify the default address family for the session, and the **set default-safi** command is used to specify the default subaddress family for the session. See the for detailed information and syntax for the **set default-afi** and **set default-safi** commands. If you do not specify a default address family, the default address family is IPv4. If you do not specify a default subaddress family, the default subaddress family is unicast.

Border Gateway Protocol (BGP) contains a separate routing table for each address family and subaddress family combination that has been configured. The address family and subaddress family options specify the

routing table to be examined. If the **all** keyword is specified for subaddress family, all subaddress family routing tables are examined.

The **show bgp cidr-only** command applies only for IPv4 prefixes. If the **ipv4** keyword is not specified and the default address family is not IPv4, the command is not available.

Use the **show bgp cidr-only** command to display CIDR routes. Routes that have their correct class (class A, B, or C) prefix length are not displayed.

Task ID

Task ID	Operations
bgp	read

Examples

The following is sample output from the **show bgp cidr-only** command in XR EXEC mode:

```
RP/0/RP0/CPU0:router# show bgp cidr-only

BGP router identifier 172.20.1.1, local AS number 1820
BGP main routing table version 2589
Dampening enabled
BGP scan interval 60 secs
Status codes: s suppressed, d damped, h history, * valid, > best
              i - internal, S stale
Origin codes: i - IGP, e - EGP, ? - incomplete
Network            Next Hop          Metric   LocPrf   Weight   Path
*> 192.0.0.0/8     192.168.72.24     0        1878     ?
*> 192.168.0.0/16  192.168.72.30     0        108      ?
```

This table describes the significant fields shown in the display.

Table 9: show bgp cidr-only Field Descriptions

Field	Description
BGP router identifier	BGP identifier for the local system.
local AS number	Autonomous system number for the local system.
BGP main routing table version	Last version of the BGP database that was installed into the main routing table.
Dampening enabled	Displayed if dampening is enabled for the routes in this BGP routing table.
BGP scan interval	Interval (in seconds) between scans of the BGP table specified by the address family and subaddress family.

Field	Description
Status codes	Status of the table entry. The status is displayed as a three-character field at the beginning of each line in the table. The first character may be (in order of precedence): S—Path is stale, indicating that a graceful restart is in progress with the peer from which the route was learned. s—Path is more specific than a locally sourced aggregate route and has been suppressed. *—Path is valid. The second character may be (in order of precedence): >—Path is the best path to use for that network. d—Path is dampened. h—Path is a history entry, representing a route that is currently withdrawn, but that is being maintained to preserve dampening information. Such routes should never be marked as valid. The third character may be: i—Path was learned by an internal BGP (iBGP) session.
Origin codes	Origin of the path. The origin code is displayed at the end of each line in the table. It can be one of the following values: i—Path originated from an Interior Gateway Protocol (IGP) and was sourced by BGP using a **network** or **aggregate-address** command. e—Entry originated from an Exterior Gateway Protocol (EGP). ?—Origin of the path is not clear. Usually, this is a route that is redistributed into BGP from an IGP.
Network	IP prefix and prefix length for a network.
Next Hop	IP address of the next system that is used when a packet is forwarded to the destination network. An entry of 0.0.0.0 indicates that the router has a non-BGP route to this network.
Metric	Value of the interautonomous system metric, otherwise known as the Multi Exit Discriminator (MED) metric.
LocPrf	Local preference value. This is used to determine the preferred exit point from the local autonomous system. It is propagated throughout the local autonomous system.
Weight	Path weight. Weight is used in choosing the preferred path to a route. It is not advertised to any neighbor.
Path	Autonomous system path to the destination network. At the end of the path is the origin code for the path.

show bgp community

To display routes that have the specified Border Gateway Protocol (BGP) communities, use the **show bgp community** command in XR EXEC mode.

show bgp [**ipv4** {**unicast** | **multicast** | **labeled-unicast** | **all** | **tunnel** | **mdt**}] **community** *community-list* [**exact-match**]
show bgp [**ipv6** {**unicast** | **multicast** | **labeled-unicast** | **all**}] **community** *community-list* [**exact-match**]
show bgp [**all** {**unicast** | **multicast** | **labeled-unicast** | **all** | **tunnel**}] **community** *community-list* [**exact-match**]
show bgp [**vpnv4 unicast** [**rd** *rd-address*]] **community** *community-list* [**exact-match**]
show bgp [**vrf** {*vrf-name* | **all**} [**ipv4** | {**unicast** | **labeled-unicast**} | **ipv6 unicast**]] **community** *community-list* [**exact-match**]
show bgp [**vpnv6 unicast** [**rd** *rd-address*]] **community** *community-list* [**exact-match**]

Syntax Description

ipv4	(Optional) Specifies IP Version 4 address prefixes.
unicast	(Optional) Specifies unicast address prefixes.
multicast	(Optional) Specifies multicast address prefixes.
labeled-unicast	(Optional) Specifies labeled unicast address prefixes.
all	(Optional) For subaddress families, specifies prefixes for all subaddress families.
tunnel	(Optional) Specifies tunnel address prefixes.
mdt	(Optional) Specifies multicast distribution tree (MDT) address prefixes.
ipv6	(Optional) Specifies IP Version 6 address prefixes.
all	(Optional) For address family, specifies prefixes for all address families.
vpnv4 unicast	(Optional) Specifies VPNv4 unicast address families.
rd *rd-address*	(Optional) Displays routes with a specific route distinguisher.
vrf	(Optional) Specifies VPN routing and forwarding (VRF) instance.
vrf-name	(Optional) Name of a VRF.
all	(Optional) For VRF, specifies all VRFs.
ipv4 { **unicast** | **labeled-unicast** }	(Optional) For VRF, specifies IPv4 unicast or labeled-unicast address families.
ipv6 unicast	(Optional) For VRF, specifies IPv6 unicast address families.
vpnv6 unicast	(Optional) Specifies VPNv6 unicast address families.
community	Specifies that only routes with communities specified by *community-list* is displayed.

community-list	Between one and seven communities. Each community can be a number in the range from 1 to 4294967295, a community specified in AA:NN format, or one of the following well-known communities: **graceful-shutdown** — Reduced preference for shutdown (well-known community) **local-AS** —Well-known community with value 4294967043 or hex 0xFFFFFF03. Routes with this community value are not advertised outside the local autonomous system or confederation boundary. **no-advertise** —Well-known community with value 4294967042 or hex 0xFFFFFF02. Routes with this community value are not advertised to any BGP peers. **no-export** —Well-known community with value 4294967041 or hex 0xFFFFFF01. Routes with this community are not advertised to external BGP peers, even if those peers are in the same confederation as the local router. **internet** —Well-known community whose value is not defined in BGP RFC. IOS XR BGP uses a value of 0 for the internet community. Routes with this community are advertised to all peers without any restrictions. For the AA:NN format: AA—Range is 0 to 65535. NN—Range is 1 to 4294967295. Up to seven community numbers can be specified.
exact-match	(Optional) Displays those routes that have communities exactly matching the specified communities.

Command Default

If no address family or subaddress family is specified, the default address family and subaddress family specified using the **set default-afi** and **set default-safi** commands are used.

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

Note The **set default-afi** command is used to specify the default address family for the session, and the **set default-safi** command is used to specify the default subaddress family for the session. See the for detailed information and syntax for the **set default-afi** and **set default-safi** commands. If you do not specify a default address family, the default address family is IPv4. If you do not specify a default subaddress family, the default subaddress family is unicast.

BGP contains a separate routing table for each configured address family and subaddress family combination. The address family and subaddress family options specify the routing table to be examined. If the **all** keyword is specified for the address family or the subaddress family, each matching routing table is examined in turn.

If more than seven communities are required, it is necessary to configure a route policy and use the show bgp route-policy, on page 421 command.

Use the **exact-match** keyword to display only those routes with a set of communities exactly matching the list of specified communities. If you omit the **exact-match** keyword, those routes containing at least the specified communities are displayed.

Task ID

Task ID	Operations
bgp	read

The following is sample output from the **show bgp community graceful-shutdown** command displaying the graceful maintenance feature information:

```
RP/0/0/CPU0:R4#show bgp community graceful-shutdown
Tue Jan 27 13:36:25.006 PST
BGP router identifier 192.168.0.4, local AS number 4
BGP generic scan interval 60 secs
BGP table state: Active
Table ID: 0xe0000000   RD version: 18
BGP main routing table version 18
BGP scan interval 60 secs

Status codes: s suppressed, d damped, h history, * valid, > best
              i - internal, r RIB-failure, S stale, N Nexthop-discard
Origin codes: i - IGP, e - EGP, ? - incomplete
   Network            Next Hop            Metric LocPrf Weight Path
*  5.5.5.5/32         10.10.10.1                88      0 1 ?

Processed 1 prefixes, 1 paths
```

Examples

The following is sample output from the **show bgp community** command in XR EXEC mode:

```
RP/0/RP0/CPU0:router# show bgp community 1820:1 exact-match

BGP router identifier 172.20.1.1, local AS number 1820
BGP main routing table version 55
Dampening enabled
BGP scan interval 60 secs
Status codes: s suppressed, d damped, h history, * valid, > best
              i - internal, S stale
Origin codes: i - IGP, e - EGP, ? - incomplete
   Network            Next Hop            Metric LocPrf Weight Path
*  10.13.0.0/16       192.168.40.24                          0 1878 704 701 200 ?
*  10.16.0.0/16       192.168.40.24                          0 1878 704 701 i
```

This table describes the significant fields shown in the display.

Table 10: show bgp community Field Descriptions

Field	Description
BGP router identifier	BGP identifier for the local system.
local AS number	Autonomous system number for the local system.

Field	Description
BGP main routing table version	Last version of the BGP database that was installed into the main routing table.
Dampening enabled	Displayed if dampening is enabled for the routes in this BGP routing table.
BGP scan interval	Interval (in seconds) between scans of the BGP table specified by the address family and subaddress family.
Status codes	Status of the table entry. The status is displayed as a three character field at the beginning of each line in the table. The first character may be (in order of precedence): S—Path is stale, indicating that a graceful restart is in progress with the peer from which the route was learned. s—Path is more specific than a locally sourced aggregate route and has been suppressed. *—Path is valid. The second character may be (in order of precedence): >—Path is the best path to use for that network. d—Path is dampened. h—Path is a history entry, representing a route that is currently withdrawn, but that is being maintained to preserve dampening information. Such routes should never be marked as valid. The third character may be: i—Path was learned by an internal BGP (iBGP) session.
Origin codes	Origin of the path. The origin code is displayed at the end of each line in the table. It can be one of the following values: i—Path originated from an Interior Gateway Protocol (IGP) and was advertised with a **network** or **aggregate-address** command. e—Path originated from an Exterior Gateway Protocol (EGP). ?—Origin of the path is not clear. Usually, this is a route that is redistributed into BGP from an IGP.
Network	IP prefix and prefix length for a network.
Next Hop	IP address of the next system that is used when a packet is forwarded to the destination network. An entry of 0.0.0.0 indicates that the router has a non-BGP route to this network.
Metric	Value of the interautonomous system metric, otherwise known as the Multi Exit Discriminator (MED) metric.
LocPrf	Local preference value. This is used to determine the preferred exit point from the local autonomous system. It is propagated throughout the local autonomous system.

Field	Description
Weight	Path weight. Weight is used in choosing the preferred path to a route. It is not advertised to any neighbor.
Path	Autonomous system path to the destination network. At the end of the path is the origin code for the path.

show bgp convergence

To display whether a specific address family has reached convergence, use the **show bgp convergence** command in XR EXEC mode.

show bgp [**ipv4** {**unicast** | **multicast** | **labeled-unicast** | **all** | **tunnel** | **mdt**}] **convergence**
show bgp [**ipv6** {**unicast** | **multicast** | **labeled-unicast** | **all**}] **convergence**
show bgp [**all** {**unicast** | **multicast** | **labeled-unicast** | **all** | **mdt** | **tunnel**}] **convergence**
show bgp [**vpnv4 unicast**] **convergence**
show bgp [**vpnv4 unicast**] **convergence**

Syntax Description

ipv4	(Optional) Specifies the IP Version 4 address family.
unicast	(Optional) Specifies the unicast address family.
multicast	(Optional) Specifies the multicast address family.
labeled-unicast	(Optional) Specifies unicast address prefixes.
all	(Optional) For subaddress family, specifies all subaddress families.
tunnel	(Optional) Specifies tunnel address prefixes.
mdt	(Optional) Specifies multicast distribution tree (MDT) address prefixes.
ipv6	(Optional) Specifies the IP Version 6 address family.
all	(Optional) For address family, specifies all address families.
vpnv4 unicast	(Optional) Specifies VPNv4 unicast address families.
vpnv6 unicast	(Optional) Specifies VPNv6 unicast address families.

Command Default

If no address family or subaddress family is specified, the default address family and subaddress family specified using the **set default-afi** and **set default-safi** commands are used.

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines



Note The **set default-afi** command is used to specify the default address family for the session, and the **set default-safi** command is used to specify the default subaddress family for the session. See the for detailed information and syntax for the **set default-afi** and **set default-safi** commands. If you do not specify a default address family, the default address family is IPv4. If you do not specify a default subaddress family, the default subaddress family is unicast.

Border Gateway Protocol (BGP) contains a separate routing table for each configured address family and subaddress family combination. The address family and subaddress family options specify the routing table to be examined. If the **all** keyword is specified for the address family or subaddress family, each matching routing table is examined in turn.

Use the **show bgp convergence** command to see if there is any pending work for BGP to perform. The software checks the following conditions to determine whether the specified address family has converged. If all the conditions are true, the address family is considered converged.

- All received updates have been processed and best routes selected.
- All selected routes have been installed in the global Routing Information Base (RIB).
- All selected routes have been advertised to peers, including any peers that are not established (unless those peers have been administratively shut down). See the **shutdown (BGP)** command for more information about administrative shutdown.

While testing that all selected routes have been advertised to peers, the **show bgp convergence**command checks the size of the write queue for each neighbor. Because this queue is shared by all address families, there is a small possibility that the command indicates the address family has not converged when, in fact, it has converged. This could happen if the neighbor write queue contained messages from some other address family.

If the specified address family has not converged, the **show bgp convergence** command output does not indicate the amount of work that is pending. To display this information, use the**show bgp summary** command.

Task ID

Task ID	Operations
bgp	read

Examples

The following shows the result of using the **show bgp convergence** command for an address family that has converged:

```
RP/0/RP0/CPU0:router# show bgp convergence

Converged.
All received routes in RIB, all neighbors updated.
All neighbors have empty write queues.
```

The following shows the result of using the **show bgp convergence** command for an address family that has not converged:

```
RP/0/RP0/CPU0:router# show bgp convergence
```

```
Not converged.
Received routes may not be entered in RIB.
One or more neighbors may need updating.
```

This table describes the significant fields shown in the display.

Table 11: show bgp convergence Field Descriptions

Field	Description
Converged/Not converged	Specifies whether or not all routes have been installed in the RIB and updates have been generated and sent to all neighbors.
[All] Received routes...	For convergence, all routes must have been installed into the RIB and all updates must have been generated. For non-convergence, some routes may not be installed in the RIB, or some routes that have been withdrawn have not yet been removed from the RIB, or some routes that are up to date in the RIB have not been advertised to all neighbors.
[All | One or more] neighbors...	Specifies the status of neighbor updating.

show bgp dampened-paths

To display Border Gateway Protocol (BGP) dampened routes, use the **show bgp dampened-paths** command in XR EXEC mode.

show bgp [**ipv4** {**unicast** | **multicast** | **labeled-unicast** | **all**}] **dampened-paths** [**standby**]
show bgp [**ipv6** {**unicast** | **multicast** | **labeled-unicast** | **all**}] **dampened-paths** [**standby**]
show bgp [**all** {**unicast** | **multicast** | **labeled-unicast** | **all** | **tunnel**}] **dampened-paths** [**standby**]
show bgp [**vpnv4 unicast** [**rd** *rd-address*]] **dampened-paths** [**standby**]
show bgp [**vrf** {*vrf-name* | **all**} [**ipv4** | {**unicast** | **labeled-unicast**} | **ipv6 unicast**]] **dampened-paths** [**standby**]
show bgp [**vpnv6 unicast** [**rd** *rd-address*]] **dampened-paths** [**standby**]

Syntax Description

ipv4	(Optional) Specifies IP Version 4 address prefixes.
unicast	(Optional) Specifies unicast address prefixes.
multicast	(Optional) Specifies multicast address prefixes.
labeled-unicast	(Optional) Specifies labeled unicast address prefixes.
all	(Optional) For subaddress families, specifies prefixes for all subaddress families.
ipv6	(Optional) Specifies IP Version 6 address prefixes.
all	(Optional) For address family, specifies prefixes for all address families.
vpnv4 unicast	(Optional) Specifies VPNv4 unicast address families.
rd *rd-address*	(Optional) Displays routes with a specific route distinguisher.
vrf	(Optional) Specifies VPN routing and forwarding (VRF) instance.
vrf-name	(Optional) Name of a VRF.
all	(Optional) For VRF, specifies all VRFs.
ipv4 {unicast | **labeled-unicast** }	(Optional) For VRF, specifies IPv4 unicast or labeled-unicast address families.
ipv6 unicast	(Optional) For VRF, specifies IPv6 unicast address families.

Command Default

If no address family or subaddress family is specified, the default address family and subaddress family specified using the **set default-afi** and **set default-safi** commands are used.

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

The **set default-afi** command is used to specify the default address family for the session, and the **set default-safi** command is used to specify the default subaddress family for the session. See the for detailed information and syntax for the **set default-afi** and **set default-safi** commands. If you do not specify a default address family, the default address family is IPv4. If you do not specify a default subaddress family, the default subaddress family is unicast.

BGP contains a separate routing table for each configured address family and subaddress family combination. The address family and subaddress family options specify the routing table to be examined. If the **all** keyword is specified for the address family or for the subaddress family, each matching routing table is examined in turn.

Task ID

Task ID	Operations
bgp	read

Examples

The following is sample output from the **show bgp dampened-paths**command in XR EXEC mode:

```
RP/0/RP0/CPU0:router# show bgp dampened-paths

BGP router identifier 10.2.0.1, local AS number 3
BGP generic scan interval 60 secs
BGP main routing table version 7
Dampening enabled
BGP scan interval 60 secs
Status codes:s suppressed, d damped, h history, * valid, > best
             i - internal, S stale

Origin codes:i - IGP, e - EGP, ? - incomplete
   Network            From            Reuse    Path
*d 10.0.0.0           10.0.101.35     00:01:20 35 i
```

This table describes the significant fields shown in the display.

Table 12: show bgp dampened-paths Field Descriptions

Field	Description
BGP router identifier	BGP identifier for the local system.
local AS number	Autonomous system number for the local system.
BGP generic scan interval	Interval (in seconds) between scans of the BGP table by a generic scanner.
BGP main routing table version	Last version of the BGP database that was installed into the main routing table.
Dampening enabled	Displayed if dampening is enabled for the routes in this BGP routing table.
BGP scan interval	Interval (in seconds) between scans of the BGP table specified by the address family and subaddress family.

Field	Description
Status codes	Status of the table entry. The status is displayed as a three-character field at the beginning of each line in the table. The first character may be (in order of precedence): S—Path is stale, indicating that a graceful restart is in progress with the peer from which the route was learned. s—Path is more specific than a locally sourced aggregate route and has been suppressed. *—Path is valid. The second character may be (in order of precedence): >—Path is the best path to use for that network. d—Path is dampened. h—Path is a history entry, representing a route that is currently withdrawn, but that is being maintained to preserve dampening information. Such routes should never be marked as valid. The third character may be: i—Path was learned by an internal BGP (iBGP) session.
Origin codes	Origin of the path. The origin code is displayed at the end of each line in the table. It can be one of the following values: i—Path originated from an Interior Gateway Protocol (IGP) and was advertised with a **network** or **aggregate-address** command. e—Path originated from an Exterior Gateway Protocol (EGP). ?—Origin of the path is not clear. Usually, this is a route that is redistributed into BGP from an IGP.
Network	IP prefix and prefix length for a network.
From	Neighbor from which the route was received.
Reuse	Time (in hours:minutes:seconds) after which the path is made available.
Path	Autonomous system path to the destination network. At the end of the path is the origin code for the path.

show bgp flap-statistics

To display information about Border Gateway Protocol (BGP) paths that have flapped, use the **show bgp flap-statistics** command in XR EXEC mode.

show bgp [**ipv4** {**unicast** | **multicast** | **labeled-unicast** | **all**}] **flap-statistics** [**regexp** *regular-expression* | **route-policy** *route-policy-name* | **cidr-only** | {*ip-address* | {*mask /prefix-length*}}] [**longer-prefixes**] [**detail**]
show bgp [**ipv6** {**unicast** | **multicast** | **labeled-unicast** | **all**}] **flap-statistics** [**regexp** *regular-expression* | **route-policy** *route-policy-name* | **cidr-only** | {*ip-address* | {*mask /prefix-length*}}] [**longer-prefixes**] [**detail**]
show bgp [**all** {**unicast** | **multicast** | **labeled-unicast** | **all**}] **flap-statistics** [**regexp** *regular-expression* | **route-policy** *route-policy-name* | **cidr-only** | {*ip-address* | {*mask /prefix-length*}}] [**longer-prefixes**] [**detail**]
show bgp [**vpnv4 unicast** [**rd** *rd-address*]] **flap-statistics** [**regexp** *regular-expression* | **route-policy** *route-policy-name* | **cidr-only** | {*ip-address* | {*mask /prefix-length*}}] [**longer-prefixes**] [**detail**]
show bgp [**vrf** {*vrf-name* | **all**} [**ipv4** | {**unicast** | **labeled-unicast**} | **ipv6 unicast**]] **flap-statistics** [**regexp** *regular-expression* | **route-policy** *route-policy-name* | **cidr-only** | {*ip-address* | {*mask /prefix-length*}}] [**longer-prefixes**] [**detail**]
show bgp [**vpnv6 unicast** [**rd** *rd-address*]] **flap-statistics** [**regexp** *regular-expression* | **route-policy** *route-policy-name* | **cidr-only** | {*ip-address* | {*mask /prefix-length*}}] [**longer-prefixes**] [**detail**]

Syntax Description

ipv4	(Optional) Specifies IP Version 4 address prefixes.
unicast	(Optional) Specifies unicast address prefixes.
multicast	(Optional) Specifies multicast address prefixes.
labeled-unicast	(Optional) Specifies labeled unicast address prefixes.
all	(Optional) For subaddress families, specifies prefixes for all subaddress families.
ipv6	(Optional) Specifies IP Version 6 address prefixes.
all	(Optional) For address family, specifies prefixes for all address families.
vpnv4 unicast	(Optional) Specifies VPNv4 unicast address families.
rd *rd-address*	(Optional) Displays routes with a specific route distinguisher.
vrf	(Optional) Specifies VPN routing and forwarding (VRF) instance.
vrf-name	(Optional) Name of a VRF.
all	(Optional) For VRF, specifies all VRFs.
ipv4 { **unicast** | **labeled-unicast** }	(Optional) For VRF, specifies IPv4 unicast or labeled-unicast address families.
ipv6 unicast	(Optional) For VRF, specifies IPv6 unicast address families.

regexp *regular-expression*	(Optional) Displays flap statistics for all paths that match the regular expression.
route-policy *route-policy-name*	(Optional) Displays flap statistics for a route policy.
cidr-only	(Optional) Displays only routes whose prefix length does not match the classful prefix length for that network. The **cidr-only** keyword can be specified only if the address family is IPv4.
ip-address	(Optional) Flap statistics for a network address only.
mask	(Optional) Network mask applied to the *ip-address* argument.
/ prefix-length	(Optional) Length of the IP address prefix. A decimal value that indicates how many of the high-order contiguous bits of the address compose the prefix (the network portion of the address). A slash (/) must precede the decimal value.
longer-prefixes	(Optional) Displays flap statistics for the specified prefix and more-specific prefixes. The **longer-prefixes** keyword is available when the *ip-address* and *mask* or */prefix-length* arguments are specified.
detail	(Optional) Displays dampening parameters for the path. The **detail** keyword cannot be specified if the **longer-prefixes** keyword is specified. The **detail** keyword is available when the *ip-address* argument or *ip-address* and *mask* or */prefix-length* arguments are specified.

Command Default

If no address family or subaddress family is specified, the default address family and subaddress family specified using the **set default-afi** and **set default-safi** commands are used.

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

Note The **set default-afi** command is used to specify the default address family for the session, and the **set default-safi** command is used to specify the default subaddress family for the session. See the for detailed information and syntax for the **set default-afi** and **set default-safi** commands. If you do not specify a default address family, the default address family is IPv4. If you do not specify a default subaddress family, the default subaddress family is unicast.

BGP contains a separate routing table for each configured address family and subaddress family combination. The address family and subaddress family options specify the routing table to be examined. If the **all** keyword is specified for the address family or subaddress family, each matching routing table is examined in turn.

Flap statistics are maintained only for paths if dampening is enabled using the **bgp dampening** command. If dampening is not enabled, the **show bgp flap-statistics** command does not display any paths.

If no arguments or keywords are specified, the software displays flap statistics for all paths for the specified address family. You can use the **regexp** , **filter-list**, **cidr-only**, and **longer-prefixes** options to limit the set of paths displayed.

If you specify a network address without a mask or prefix length, the longest matching prefix for the network address is displayed. When displaying flap statistics for a single route, use the **detail** keyword to display dampening parameters for the route.

Task ID

Task ID	Operations
bgp	read

Examples

The following is sample output from the **show bgp flap-statistics** command:

```
RP/0/RP0/CPU0:router# show bgp flap-statistics

BGP router identifier 172.20.1.1, local AS number 1820
BGP main routing table version 26180
Dampening enabled
BGP scan interval 60 secs
Status codes: s suppressed, d damped, h history, * valid, > best
              i - internal, S stale
Origin codes: i - IGP, e - EGP, ? - incomplete
Network            From           Flaps Duration Reuse    Path
*d 10.0.0.0        172.20.16.177  4     00:13:31 00:18:10 100
*d 10.10.0.0       172.20.16.177  4     00:02:45 00:28:20 100
```

The following is sample output from the **show bgp flap-statistics** command with the **detail** keyword in XR EXEC mode:

```
RP/0/RP0/CPU0:router# show bgp flap-statistics 172.31.12.166 detail

BGP router identifier 10.0.0.5, local AS number 1
BGP main routing table version 738
Dampening enabled
BGP scan interval 60 secs
Status codes: s suppressed, d damped, h history, * valid, > best
              i - internal, S stale
Origin codes: i - IGP, e - EGP, ? - incomplete
   Network            From            Flaps Duration Reuse    Path
 h 172.31.12.166            10.0.101.1      6     00:03:28          2 2000 3000

   Half life        Suppress        Reuse penalty    Max. supp. time
   00:15:00         2000            750              01:00:00
```

This table describes the significant fields shown in the display.

Table 13: show bgp flap-statistics Field Descriptions

Field	Description
BGP route identifier	BGP identifier for the local system.
local AS number	Autonomous system number for the local system.
BGP main routing table version	Last version of the BGP database that was installed into the main routing table.
Dampening enabled	Displayed if dampening has been enabled for the routes in this BGP routing table.
BGP scan interval	Interval (in seconds) between scans of the BGP table specified by the address family and subaddress family.
Status codes	Status of the table entry. The status is displayed as a three-character field at the beginning of each line in the table. The first character may be (in order of precedence): S—Path is stale, indicating that a graceful restart is in progress with the peer from which the route was learned. s—Path is more specific than a locally sourced aggregate route and has been suppressed. *—Path is valid. The second character may be (in order of precedence): >—Path is the best path to use for that network. d—Path is dampened. h—Path is a history entry, representing a route that is currently withdrawn, but that is being maintained to preserve dampening information. Such routes should never be marked as valid. The third character may be: i—Path was learned by an internal BGP (iBGP) session.
Origin codes	Origin of the path. The origin code is displayed at the end of each line in the table. It can be one of the following values: i—Path originated from an Interior Gateway Protocol (IGP) and was advertised with a network or aggregate-address command. e—Path originated from an Exterior Gateway Protocol (EGP). ?—Origin of the path is not clear. Usually, this is a route that is redistributed into BGP from an IGP.
Network	IP prefix and prefix length for a network that is dampened.
From	IP address of the peer that advertised this route.
Flaps	Number of times the route has flapped.

Field	Description
Duration	Time (in hours:minutes:seconds) since the first flap.
Reuse	Time (in hours:minutes:seconds) after which the path is made available.
Path	Autonomous system path of the route that is being dampened.
Half life	Half-life value used when dampening this route. The half-life is the amount of time that must elapse to reduce the reuse penalty by half. The half-life value is specified using the **bgp dampening** command.
Suppress	Suppress value used to dampen this route. The suppress value is the value that the penalty must exceed for the route to be suppressed. The suppress value can be configured using the **bgp dampening** command.
Reuse penalty	Reuse penalty used to dampen this route. The penalty must fall below the reuse penalty for the route to be unsuppressed. The reuse penalty can be configured using the **bgp dampening** command.
Max supp. time	Maximum length of time that the route may be suppressed due to dampening. The maximum suppress time can be configured using the **bgp dampening** command.

show bgp inconsistent-as

To display Border Gateway Protocol (BGP) routes originated from more than one autonomous system, use the **show bgp inconsistent-as** command in XR EXEC mode.

show bgp [**ipv4** {**unicast** | **multicast** | **labeled-unicast** | **all** | **tunnel** | **mdt**}] **inconsistent-as** [**standby**]
show bgp [**ipv6** {**unicast** | **multicast** | **labeled-unicast** | **all**}] **inconsistent-as** [**standby**]
show bgp [**all** {**unicast** | **multicast** | **labeled-unicast** | **all** | **tunnel** | **mdt**}] **inconsistent-as** [**standby**]
show bgp vpnv4 unicast [**rd** *rd-address*] **inconsistent-as** [**standby**]
show bgp [**vrf** {*vrf-name* | **all**} [**ipv4** | {**unicast** | **labeled-unicast**} | **ipv6 unicast**]] **inconsistent-as** [**standby**]
show bgp [**vpnv6 unicast** [**rd** *rd-address*]] **inconsistent-as** [**standby**]

Syntax Description

ipv4	(Optional) Specifies IP Version 4 address prefixes.
unicast	(Optional) Specifies unicast address prefixes.
multicast	(Optional) Specifies multicast address prefixes.
labeled-unicast	(Optional) Specifies labeled unicast address prefixes.
all	(Optional) For subaddress families, specifies prefixes for all subaddress families.
tunnel	(Optional) Specifies tunnel address prefixes.
ipv6	(Optional) Specifies IP Version 6 address prefixes.
all	(Optional) For address family, specifies prefixes for all address families.
vpnv4 unicast	(Optional) Specifies VPNv4 unicast address families.
rd *rd-address*	(Optional) Displays routes with a specific route distinguisher.
vrf	(Optional) Specifies VPN routing and forwarding (VRF) instance.
vrf-name	(Optional) Name of a VRF.
all	(Optional) For VRF, specifies all VRFs.
ipv4 { **unicast** | **labeled-unicast** }	(Optional) For VRF, specifies IPv4 unicast or labeled-unicast address families.
ipv6 unicast	(Optional) For VRF, specifies IPv6 unicast address families.

Command Default

If no address family or subaddress family is specified, the default address family and subaddress family specified using the **set default-afi** and **set default-safi** commands are used.

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

Note The **set default-afi** command is used to specify the default address family for the session, and the **set default-safi** command is used to specify the default subaddress family for the session. See the for detailed information and syntax for the **set default-afi** and **set default-safi** commands. If you do not specify a default address family, the default address family is IPv4. If you do not specify a default subaddress family, the default subaddress family is unicast.

BGP contains a separate routing table for each configured address family and subaddress family combination. The address family and subaddress family options specify the routing table to be examined. If the **all** keyword is specified for the address family or for the subaddress family, each matching routing table is examined in turn.

Use the **show bgp inconsistent-as**command to search through all prefixes in the specified BGP routing table and display the paths for any prefix that has inconsistent originating autonomous system numbers. The originating autonomous system is the last autonomous system number displayed in the path field and should be the same for all paths.

If a prefix has one or more paths originating from different autonomous systems, all paths for that prefix are displayed.

Task ID

Task ID	Operations
bgp	read

Examples

The following is sample output from the **show bgp inconsistent-as** command in XR EXEC mode:

```
RP/0/RP0/CPU0:router# show bgp inconsistent-as

BGP router identifier 172.20.1.1, local AS number 1820
BGP main routing table version 1129
BGP scan interval 60 secs
Status codes: s suppressed, d damped, h history, * valid, > best
              i - internal, S stale
Origin codes: i - IGP, e - EGP, ? - incomplete
Network            Next Hop          Metric             LocPrf Weight Path
*  10.0.0.0        172.16.232.55          0             0 300 88 90 99 ?
*>                 172.16.232.52       2222             0 400 ?
*  172.16.0.0      172.16.232.55          0             0 300 90 99 88 200 ?
*>                 172.16.232.52       2222             0 400 ?
* 192.168.199.0    172.16.232.55          0             0 300 88 90 99 ?
*>                 172.16.232.52       2222             0 400 ?
```

This table describes the significant fields shown in the display.

Table 14: show bgp inconsistent-as Field Descriptions

Field	Description
BGP router identifier	BGP identifier for the local system.
local AS number	Autonomous system number for the local system.
BGP main routing table version	Last version of the BGP database that was installed into the main routing table.
Dampening enabled	Displayed if dampening is enabled for the routes in this BGP routing table.
BGP scan interval	Interval (in seconds) between scans of the BGP table specified by the address family and subaddress family.
Status codes	Status of the table entry. The status is displayed as a three-character field at the beginning of each line in the table. The first character may be (in order of precedence): S—Path is stale, indicating that a graceful restart is in progress with the peer from which the route was learned. s—Path is more specific than a locally sourced aggregate route and has been suppressed. *—Path is valid. The second character may be (in order of precedence): >—Path is the best path to use for that network. d—Path is dampened. h—Path is a history entry, representing a route that is currently withdrawn, but that is being maintained to preserve dampening information. Such routes should never be marked as valid. The third character may be: i—Path was learned by an internal BGP (iBGP) session.
Origin codes	Origin of the path. The origin code is displayed at the end of each line in the table. It can be one of the following values: i—Path originated from an Interior Gateway Protocol (IGP) and was advertised with a **network** or **aggregate-address** command. e—Path originated from an Exterior Gateway Protocol (EGP). ?—Origin of the path is not clear. Usually, this is a route that is redistributed into BGP from an IGP.
Network	IP prefix and prefix length for a network.

Field	Description
Next Hop	IP address of the next system that is used when a packet is forwarded to the destination network. An entry of 0.0.0.0 indicates that the router has a non-BGP route to this network.
Metric	Value of the interautonomous system metric, otherwise known as the Multi Exit Discriminator (MED) metric.
LocPrf	Local preference value. This is used to determine the preferred exit point from the local autonomous system. It is propagated throughout the local autonomous system.
Weight	Path weight. Weight is used in choosing the preferred path to a route. It is not advertised to any neighbor.
Path	Autonomous system path to the destination network. At the end of the path is the origin code for the path.

show bgp labels

To display Border Gateway Protocol (BGP) routes and their incoming and outgoing labels, use the **show bgp labels**command in XR EXEC mode.

show bgp labels

Syntax Description

ipv4	(Optional) Specifies IP Version 4 address prefixes.
unicast	(Optional) Specifies unicast address prefixes.
multicast	(Optional) Specifies multicast address prefixes.
labeled-unicast	(Optional) Specifies labeled-unicast address prefixes.
all	(Optional) For subaddress families, specifies prefixes for all subaddress families.
ipv6	(Optional) Specifies IP Version 6 address prefixes.
vpnv4 unicast	(Optional) Specifies VPNv4 unicast address families.
rd *rd-address*	(Optional) Displays routes with a specific route distinguisher.
vrf	(Optional) Specifies VPN routing and forwarding (VRF) instance.
vrf-name	(Optional) Name of a VRF.
all	(Optional) For VRF, specifies all VRFs.
ipv4 {unicast | **labeled-unicast** }	(Optional) For VRF, specifies IPv4 unicast or labeled-unicast address families.
ipv6 unicast	(Optional) For VRF, specifies IPv6 unicast address families.

Command Default

If no address family or subaddress family is specified, the default address family and subaddress family specified using the **set default-afi** and **set default-safi** commands are used.

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

Task ID

Task ID	Operations
bgp	read

Examples

The following is sample output from the **show bgp labels**command in XR EXEC mode:

```
RP/0/RP0/CPU0:router# show bgp vrf BAR ipv4 unicast labels

BGP VRF BAR, state: Active BGP Route Distinguisher: 100:1 BGP router identifier 10.1.1.1,
local AS number 100
BGP table state: Active BGP main routing table version 12

Status codes: s suppressed, d damped, h history, * valid, > best
              i - internal, S stale
Origin codes: i - IGP, e - EGP, ? - incomplete
   Network            Next Hop        Rcvd Label      Local Label
Route Distinguisher: 100:1 (default for vrf BAR)
*> 20.1.1.1/32        10.0.101.1      16              nolabel
*> 20.1.1.2/32        10.0.101.1      16              nolabel
*> 20.1.1.3/32        10.0.101.1      16              nolabel
*> 20.1.1.4/32        10.0.101.1      16              nolabel
*> 20.1.1.5/32        10.0.101.1      16              nolabel

Processed 5 prefixes, 5 paths
```

This table describes the significant fields shown in the display.

Table 15: show bgp labels Field Descriptions

Field	Description
BGP Route Distinguisher	BGP route distinguisher.
BGP router identifier	BGP identifier for the local system.
local AS number	Autonomous system number for the local system.
BGP table state	State of the BGP routing table.
BGP main routing table version	Last version of the BGP database that was installed into the main routing table.

Field	Description
Status codes	Status of the table entry. The status is displayed as a three-character field at the beginning of each line in the table. The first character may be (in order of precedence): S—Path is stale, indicating that a graceful restart is in progress with the peer from which the route was learned. s—Path is more specific than a locally sourced aggregate route and has been suppressed. *—Path is valid. The second character may be (in order of precedence): >—Path is the best path to use for that network. d—Path is dampened. h—Path is a history entry, representing a route that is currently withdrawn, but that is being maintained to preserve dampening information. Such routes should never be marked as valid. The third character may be: i—Path was learned by an internal BGP (iBGP) session.
Origin codes	Origin of the path. The origin code is displayed at the end of each line in the table. It can be one of the following values: i—Path originated from an Interior Gateway Protocol (IGP) and was sourced by BGP using a network or **aggregate-address** command. e—Path originated from an Exterior Gateway Protocol (EGP). ?—Origin of the path is not clear. Usually, this is a route that is redistributed into BGP from an IGP.
Network	IP prefix and prefix length for a network.
Next Hop	IP address of the next system that is used when a packet is forwarded to the destination network. An entry of 0.0.0.0 indicates that the router has a non-BGP route to this network.
Rcvd Label	Received label.
Local Label	Local label.

show bgp large-community

To display routes that have the specified Border Gateway Protocol (BGP) large-communities, use the **show bgp large-community** command in XR EXEC mode mode.

show bgp [**ipv4** { **unicast** | **multicast** | **labeled-unicast** | **all** | **tunnel** | **mdt** }] **large-community** *large-community-list* [**exact-match**]
show bgp [**ipv6** { **unicast** | **multicast** | **labeled-unicast** | **all** }] **large-community** *large-community-list* [**exact-match**]
show bgp [**all** { **unicast** | **multicast** | **labeled-unicast** | **all** | **tunnel** }] **large-community** *large-community-list* [**exact-match**]
show bgp [**vpnv4 unicast** [**rd** *rd-address*]] **large-community** *large-community-list* [**exact-match**]
show bgp [**vrf** { *vrf-name* | **all** } [**ipv4** | { **unicast** | **labeled-unicast** } | **ipv6 unicast**]] **large-community** *large-community-list* [**exact-match**]
show bgp [**vpnv6 unicast** [**rd** *rd-address*]] **large-community** *large-community-list* [**exact-match**]

Syntax Description

ipv4	(Optional) Specifies IP Version 4 address prefixes.
unicast	(Optional) Specifies unicast address prefixes.
multicast	(Optional) Specifies multicast address prefixes.
labeled-unicast	(Optional) Specifies labeled unicast address prefixes.
all	(Optional) For subaddress families, specifies prefixes for all subaddress families.
tunnel	(Optional) Specifies tunnel address prefixes.
mdt	(Optional) Specifies multicast distribution tree (MDT) address prefixes.
ipv6	(Optional) Specifies IP Version 6 address prefixes.
all	(Optional) For address family, specifies prefixes for all address families.
vpnv4 unicast	(Optional) Specifies VPNv4 unicast address families.
rd *rd-address*	(Optional) Displays routes with a specific route distinguisher.
vrf	(Optional) Specifies VPN routing and forwarding (VRF) instance.
vrf-name	(Optional) Name of a VRF.
all	(Optional) For VRF, specifies all VRFs.
ipv4 { **unicast** | **labeled-unicast** }	(Optional) For VRF, specifies IPv4 unicast or labeled-unicast address families.
ipv6 unicast	(Optional) For VRF, specifies IPv6 unicast address families.

vpnv6 unicast	(Optional) Specifies VPNv6 unicast address families.
large-community	Indicates that only routes with large communities specified by *large-community-list* is displayed.
large-community-list	They are specified as three non negative decimal integers separated by colons. For example, 1:2:3. Each integer is stored in 32 bits. The possible range for each integer is 0 to 4294967295.
exact-match	(Optional) Displays those routes that have communities exactly matching the specified communities.

Command Default

None.

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 6.3.1	This command was introduced.

Usage Guidelines

Use the **exact-match** keyword to display only those routes with a set of large communities exactly matching the list of specified large communities. If you omit the **exact-match** keyword, those routes containing at least the specified large communities are displayed.

Task ID

Task ID	Operations
bgp	read

Examples

The following is sample output from the **show bgp large-community** command in XR EXEC mode mode:

```
RP/0/RP0/CPU0:router#show bgp large-community 1:2:3 5:6:7
Thu Mar 23 14:40:33.597 PDT
BGP router identifier 4.4.4.4, local AS number 3
BGP generic scan interval 60 secs
Non-stop routing is enabled
BGP table state: Active
Table ID: 0xe0000000   RD version: 66
BGP main routing table version 66
BGP NSR Initial initsync version 3 (Reached)
BGP NSR/ISSU Sync-Group versions 66/0
BGP scan interval 60 secs

Status codes: s suppressed, d damped, h history, * valid, > best
              i - internal, r RIB-failure, S stale, N Nexthop-discard
Origin codes: i - IGP, e - EGP, ? - incomplete
   Network            Next Hop            Metric LocPrf Weight Path
*>i3.3.3.3/32         10.10.10.3               0     94      0 ?
*> 5.5.5.5/32         11.11.11.5               0             0 5 ?
```

In the above example, the routes 3.3.3.3/32 and 5.5.5.5/32 contain both the large communities 1:2:3 and 5:6:7 and some additional large communities.

```
RP/0/RP0/CPU0:router#show bgp large-community 1:2:3 5:6:7 exact-match
Thu Mar 23 14:50:28.155 PDT
```

In the above example, no routes are displayed because the available routes have some additional large-communities, so the match is not exact.

show bgp l2vpn

To display BGP routes associated with VPLS or VPWS or EVPN under L2VPN address family, use the**show bgp l2vpn** command in XR EXEC mode.

show bgp l2vpn { **vpls** | **vpws** | **evpn** } **rd** *rd _value bgp_prefix* [**detail**]

Syntax Description

vpls	Specifies Virtual Private LAN Services (VPLS).
vpws	Specifies Virtual Private Wire Service (VPWS).
evpn	Specifies Ethernet Virtual Private Network (EVPN).
rd _value	Value of the route distinguisher.
bgp_prefix	Specifies BGP prefix.
detail	Provides detailed output for the specified route distinguisher and BGP prefix.

Command Default

No default behavior or values

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

Task ID

Task ID	Operation
bgp	read

The following example is sample output from the **show bgp l2vpn vpls** for route distinguisher 1:1:

```
RP/0/RP0/CPU0:router#show bgp l2vpn vpls rd 1:1 2:1
BGP routing table entry for 2:1/32, Route Distinguisher: 1:1
Versions:
  Process           bRIB/RIB  SendTblVer
  Speaker                 2           2
    Local Label: 16000
Paths: (1 available, best #1)
  Advertised to update-groups (with more than one peer):
    0.1
  Advertised to peers (in unique update groups):
    100.100.100.1
  Path #1: Received by speaker 0
  Local
    0.0.0.0 from 0.0.0.0 (200.200.200.1)
      Origin IGP, localpref 100, valid, redistributed, best, import-candidate
      Extended community: RT:4:4 L2VPN:19:0:1500
      Block Size:10
```

The following example is sample output from the **show bgp l2vpn vpws** for route distinguisher 200:200:

```
RP/0/RP0/CPU0:router#show bgp l2vpn vpws rd 200:200 3:1
BGP routing table entry for 3:1/32, Route Distinguisher: 200:200
Versions:
  Process           bRIB/RIB  SendTblVer
  Speaker                 6           6
    Local Label: 16015
Paths: (1 available, best #1)
  Advertised to update-groups (with more than one peer):
    0.1
  Advertised to peers (in unique update groups):
    100.100.100.1
  Path #1: Received by speaker 0
  Local
    0.0.0.0 from 0.0.0.0 (200.200.200.1)
      Origin IGP, localpref 100, valid, redistributed, best, import-candidate
      Extended community: RT:2:2 L2VPN:4:0:1500
      Circuit Vector:0xfd 0xff
      Block Size:10
```

The following example is sample output from the **show bgp l2vpn vpls** for local NLRI:

RD is 3.3.3.3:3276, NH Address is 100.0.0.1, and VPLS ID is 150:200. The RT is 200:100.

Note The RT and VPLS-ID are always same for the same VPLS instance.

```
RP/0/RP0/CPU0:router#show bgp l2vpn vpls
Sat Jun  6 17:01:18.610 PST
BGP router identifier 3.3.3.3, local AS number 101
BGP generic scan interval 60 secs
BGP table state: Active
Table ID: 0x0
BGP main routing table version 5
BGP scan interval 60 secs

Status codes: s suppressed, d damped, h history, * valid, > best
              i - internal, S stale
Origin codes: i - IGP, e - EGP, ? - incomplete
   Network            Next Hop
Route Distinguisher: 3.3.3.3:3276 (default for vrf g1:b1)
*>i200.0.0.1          30.0.0.2
*>i100.0.0.1          0.0.0.0
Route Distinguisher: 2.2.2.2:3435
*>i200.0.0.1          30.0.0.2

Processed 3 prefixes, 3 paths


RP/0/RP0/CPU0:router#show bgp l2vpn vpls rd 3.3.3.3:3276  100.0.0.1
Sat Jun  6 16:40:03.191 PST
BGP routing table entry for 100.0.0.1, Route Distinguisher: 3.3.3.3:3276
Versions:
  Process           bRIB/RIB  SendTblVer
  Speaker                 3           3
    Last Modified: Jun  6 11:20:57.944 for 05:19:05
Paths: (1 available, best #1)
  Advertised to peers (in unique update groups):
    30.0.0.2
```

```
  Path #1: Received by speaker 0
  Local
    0.0.0.0 from 0.0.0.0 (3.3.3.3)
      Origin IGP, localpref 100, valid, redistributed, best, import-candidate
      Extended community: RT:200:100 VPLS-ID:150:200
```

The following example is sample output from the show bgp l2vpn vpls for remote NLRI:

RD is 2.2.2.2:3435, NH Address is 200.0.0.1, and VPLS ID is 150:200. The RT is 200:100.

```
RP/0/RP0/CPU0:router#show bgp l2vpn vpls rd 2.2.2.2:3435 200.0.0.1
Sat Jun  6 16:53:55.726 PST
BGP routing table entry for 200.0.0.1, Route Distinguisher: 2.2.2.2:3435
Versions:
  Process           bRIB/RIB  SendTblVer
  Speaker                 5           5
Last Modified: Jun  6 11:20:57.944 for 05:32:58
Paths: (1 available, best #1)
  Not advertised to any peer
  Path #1: Received by speaker 0
  Local
    30.0.0.2 from 30.0.0.2 (133.133.133.133)
      Origin IGP, localpref 100, valid, internal, best, import-candidate, imported
      Extended community: RT:200:50 VPLS-ID:150:200
```

The following example is sample output of the **show bgp l2vpn evpn** command for a Data Center Interconnect Layer 3 Gateway.

```
RP/0/RP0/CPU0:router#show bgp l2vpn evpn
Fri Aug 21 00:24:10.773 PDT
BGP router identifier 30.30.30.30, local AS number 100
BGP generic scan interval 60 secs
Non-stop routing is enabled
BGP table state: Active
Table ID: 0x0   RD version: 0
BGP main routing table version 16
BGP NSR Initial initsync version 1 (Reached)
BGP NSR/ISSU Sync-Group versions 16/0
BGP scan interval 60 secs

Status codes: s suppressed, d damped, h history, * valid, > best
              i - internal, r RIB-failure, S stale, N Nexthop-discard
Origin codes: i - IGP, e - EGP, ? - incomplete
   Network            Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 100:1
*>i[2][10000][48][0226.51bd.c81c][32][200::1001]/232
                      11.0.0.1                 100      0 i
*>i[2][10000][48][0226.51bd.c81c][32][200:1::1001]/232
                      11.0.0.1                 100      0 i
*>i[2][10000][48][0226.51bd.c81c][32][200.1.1.1]/136
                      11.0.0.1                 100      0 i
*>i[2][10000][48][0226.51bd.c81c][32][200.1.1.2]/136
                      11.0.0.1                 100      0 i
*>i[5][4231][32][100.1.1.1]/80
                      11.0.0.1                 100      0 i
*>i[5][4231][32][100.1.1.2]/80
                      11.0.0.1                 100      0 i
*>i[5][4231][112][fec0::1001]/176
                      11.0.0.1                 100      0 i
*>i[5][4232][112][fec0::1:1001]/176
                      11.0.0.1                 100      0 i

Processed 8 prefixes, 8 paths
```

The following example is sample output of the **show bgp l2vpn evpn rd** command for a Data Center Interconnect Layer 3 Gateway. This sample output provides details for the specified route distinguisher and prefix.

```
RP/0/RP0/CPU0:router# show bgp l2vpn evpn rd 100:1 [5][4231][112][fec0::1001]/176 detail
Fri Aug 21 00:34:43.747 PDT
BGP routing table entry for [5][4231][112][fec0::1001]/176, Route Distinguisher: 100:1
Versions:
  Process           bRIB/RIB  SendTblVer
  Speaker                 5           5
    Flags: 0x04040001+0x00000000;
Last Modified: Aug 21 00:16:58.000 for 00:17:46
Paths: (1 available, best #1)
  Not advertised to any peer
  Path #1: Received by speaker 0
  Flags: 0x4000600025060005, import: 0x3f
  Not advertised to any peer
  Local
    11.0.0.1 (metric 2) from 20.0.0.1 (11.0.0.1)
      Received Label 16001
      Origin IGP, localpref 100, valid, internal, best, group-best, import-candidate,
reoriginate, not-in-vrf
      Received Path ID 0, Local Path ID 1, version 5
     Extended community: Flags 0x2: Encapsulation Type:8 Router MAC:aabb.ccdd.eeff RT:65540:1
 RT:40.40.40.40:1 RT:100:1
      Originator: 11.0.0.1, Cluster list: 20.20.20.20
      EVPN ESI: ffff.ffff.ffff.ffff.ff01, Gateway Address : fec0::254
```

show bgp l2vpn vpls

To display L2VPN information on BGP summary, routes for a specified bridge group domain, advertised routes, routes with a specific route distinguisher, BGP neighbor connections, nexthops, and BGP process, use the **show bgp l2vpn vpls** command in XR EXEC mode.

show bgp l2vpn vpls {**summary** | **rd** | **neighbors** | **nexthops** | **bdomain** | **advertised** | **process**}

Syntax Description

summary	Displays the summary of BGP neighbor status.
rd	Displays routes with a specific route distinguisher.
neighbors	Displays detailed information on TCP and BGP neighbor connections.
nexthops	Shows nexthop related information.
bdomain	Displays routes for a specified Bridge Group:domain
advertised	Shows advertised routes.
process	Displays BGP process information.

Command Default

No default behavior or values

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

Task ID

Task ID	Operation
bgp	read

Example

The following example is sample output from the **show bgp l2vpn vpls** command:

```
RP/0/RP0/CPU0:router#show bgp l2vpn vpls
Wed Mar 17 15:26:29.433 EDT
BGP router identifier 60.60.60.60, local AS number 1
BGP generic scan interval 60 secs
Non-stop routing is enabled
BGP table state: Active
Table ID: 0x0
BGP main routing table version 24001
BGP NSR Initial initsync version 1 (Reached)
```

```
BGP scan interval 60 secs

Status codes: s suppressed, d damped, h history, * valid, > best
              i - internal, S stale
Origin codes: i - IGP, e - EGP, ? - incomplete
   Network            Next Hop        Rcvd Label      Local Label
Route Distinguisher: 101:1 (default for vrf bg1:bg1_bd1)
*>i10.10.10.10/32     10.10.10.10     nolabel         nolabel
*> 60.60.60.60/32     0.0.0.0         nolabel         nolabel
Route Distinguisher: 102:1 (default for vrf bg1:bg1_bd2)
*>i10.10.10.10/32     10.10.10.10     nolabel         nolabel
*> 60.60.60.60/32     0.0.0.0         nolabel         nolabel
```

The following example is sample output from the **show bgp l2vpn vpls** command with the summary keyword:

```
RP/0/RP0/CPU0:router#show bgp l2vpn vpls summary
Wed Mar 17 15:27:09.502 EDT
BGP router identifier 60.60.60.60, local AS number 1
BGP generic scan interval 60 secs
Non-stop routing is enabled
BGP table state: Active
Table ID: 0x0
BGP main routing table version 24001
BGP NSR Initial initsync version 1 (Reached)
BGP scan interval 60 secs

BGP is operating in STANDALONE mode.


Process       RcvTblVer   bRIB/RIB   LabelVer  ImportVer  SendTblVer  StandbyVer
Speaker           24001      24001      24001      24001       24001           0

Neighbor        Spk    AS MsgRcvd MsgSent   TblVer  InQ OutQ  Up/Down  St/PfxRcd
10.10.10.10       0     1   45532    8392    24001    0    0 03:06:25        8000
```

The following example is sample output from the **show bgp l2vpn vpls** command for Route Distinguisher: 101:1:

```
RP/0/RP0/CPU0:router#show bgp l2vpn vpls rd 101:1
Wed Mar 17 15:27:31.347 EDT
BGP router identifier 60.60.60.60, local AS number 1
BGP generic scan interval 60 secs
Non-stop routing is enabled
BGP table state: Active
Table ID: 0x0
BGP main routing table version 24001
BGP NSR Initial initsync version 1 (Reached)
BGP scan interval 60 secs

Status codes: s suppressed, d damped, h history, * valid, > best
              i - internal, S stale
Origin codes: i - IGP, e - EGP, ? - incomplete
   Network            Next Hop        Rcvd Label      Local Label
Route Distinguisher: 101:1 (default for vrf bg1:bg1_bd1)
*>i10.10.10.10/32     10.10.10.10     nolabel         nolabel
*> 60.60.60.60/32     0.0.0.0         nolabel         nolabel

Processed 2 prefixes, 2 paths
```

The following example is sample output from the **show bgp l2vpn vpls** command for BGP neighbor 10.10.10.10:

```
RP/0/RP0/CPU0:router#show bgp l2vpn vpls neighbors 10.10.10.10
Wed Mar 17 15:28:28.766 EDT

BGP neighbor is 10.10.10.10
 Remote AS 1, local AS 1, internal link
 Remote router ID 10.10.10.10
  BGP state = Established, up for 03:07:44
  NSR State: None
  Last read 00:00:31, Last read before reset 00:00:00
  Hold time is 180, keepalive interval is 60 seconds
  Configured hold time: 180, keepalive: 60, min acceptable hold time: 3
  Last write 00:00:45, attempted 19, written 19
  Second last write 00:01:45, attempted 19, written 19
  Last write before reset 00:00:00, attempted 0, written 0
  Second last write before reset 00:00:00, attempted 0, written 0
  Last write pulse rcvd  Mar 17 15:27:57.362 last full not set pulse count 847
  Last write pulse rcvd before reset 00:00:00
  Socket not armed for io, armed for read, armed for write
  Last write thread event before reset 00:00:00, second last 00:00:00
  Last KA expiry before reset 00:00:00, second last 00:00:00
  Last KA error before reset 00:00:00, KA not sent 00:00:00
  Last KA start before reset 00:00:00, second last 00:00:00
  Precedence: internet
  Non-stop routing is enabled
  Graceful restart is enabled
  Restart time is 300 seconds
  Stale path timeout time is 1200 seconds
  Neighbor capabilities:
    Route refresh: advertised and received
    Graceful Restart (GR Awareness): received
    4-byte AS: advertised and received
    Address family IPv4 Unicast: advertised and received
    Address family VPNv4 Unicast: advertised and received
    Address family L2VPN VPLS: advertised and received
  Received 45533 messages, 0 notifications, 0 in queue
  Sent 8393 messages, 0 notifications, 0 in queue
  Minimum time between advertisement runs is 0 secs

 For Address Family: IPv4 Unicast
  BGP neighbor version 1
  Update group: 0.2
  AF-dependant capabilities:
    Graceful Restart Capability advertised and received
      Local restart time is 300, RIB purge time is 900 seconds
      Maximum stalepath time is 1200 seconds
      Remote Restart time is 300 seconds
  Route refresh request: received 0, sent 0
  0 accepted prefixes, 0 are bestpaths
  Cumulative no. of prefixes denied: 0.
  Prefix advertised 0, suppressed 0, withdrawn 0
  Maximum prefixes allowed 524288
  Threshold for warning message 75%, restart interval 0 min
  An EoR was received during read-only mode
  Last ack version 1, Last synced ack version 0
  Outstanding version objects: current 0, max 0

 For Address Family: VPNv4 Unicast
  BGP neighbor version 1
  Update group: 0.2
  AF-dependant capabilities:
```

```
    Graceful Restart Capability advertised and received
      Local restart time is 300, RIB purge time is 900 seconds
      Maximum stalepath time is 1200 seconds
      Remote Restart time is 300 seconds
  Route refresh request: received 0, sent 0
  0 accepted prefixes, 0 are bestpaths
  Cumulative no. of prefixes denied: 0.
  Prefix advertised 0, suppressed 0, withdrawn 0
  Maximum prefixes allowed 524288
  Threshold for warning message 75%, restart interval 0 min
  An EoR was received during read-only mode
  Last ack version 1, Last synced ack version 0
  Outstanding version objects: current 0, max 0

 For Address Family: L2VPN VPLS
  BGP neighbor version 24001
  Update group: 0.2
  AF-dependant capabilities:
    Graceful Restart Capability advertised and received
      Local restart time is 300, RIB purge time is 900 seconds
      Maximum stalepath time is 1200 seconds
      Remote Restart time is 300 seconds
  Route refresh request: received 0, sent 203
  8000 accepted prefixes, 8000 are bestpaths
  Cumulative no. of prefixes denied: 18172.
    No policy: 0, Failed RT match: 18172
    By ORF policy: 0, By policy: 0
  Prefix advertised 8000, suppressed 0, withdrawn 0
  Maximum prefixes allowed 524288
  Threshold for warning message 75%, restart interval 0 min
  An EoR was received during read-only mode
  Last ack version 24001, Last synced ack version 0
  Outstanding version objects: current 0, max 2

  Connections established 1; dropped 0
  Local host: 60.60.60.60, Local port: 179
  Foreign host: 10.10.10.10, Foreign port: 50472
  Last reset 00:00:00
```

The following example is sample output from the **show bgp l2vpn vpls** command with the advertised keyword:

```
RP/0/RP0/CPU0:router#show bgp l2vpn vpls advertised
Wed Mar 17 15:29:13.787 EDT
Route Distinguisher: 101:1
60.60.60.60/32 is advertised to 10.10.10.10
  Path info:
    neighbor: Local           neighbor router id: 60.60.60.60
    valid  redistributed  best  import-candidate
  Attributes after inbound policy was applied:
    next hop: 0.0.0.0
    EXTCOMM
    origin: IGP
    aspath:
    extended community: RT:101:1 L2VPN AGI:1:101
```

The following example is sample output from the **show bgp l2vpn vpls** command with the nexthops keyword:

```
RP/0/RP0/CPU0:router#show bgp l2vpn vpls nexthops
Wed Mar 17 15:29:36.357 EDT
```

```
Total Nexthop Processing
  Time Spent: 0.000 secs

Maximum Nexthop Processing
  Received: 82y46w
  Bestpaths Deleted: 0
  Bestpaths Changed: 0
  Time Spent: 0.000 secs

Last Notification Processing
  Received: 03:10:50
  Time Spent: 0.000 secs

Gateway Address Family: IPv4 Unicast
Table ID: 0xe0000000
Nexthop Count: 2
Critical Trigger Delay: 3000msec
Non-critical Trigger Delay: 10000msec

Nexthop Version: 1, RIB version: 1

Status codes: R/UR Reachable/Unreachable
              C/NC Connected/Not-connected
              L/NL Local/Non-local
              I    Invalid (Policy Match Failed)
Next Hop        Status          Metric        Notf      LastRIBEvent RefCount
10.10.10.10     [R][NC][NL]          2         1/0    03:10:50 (Cri)     8000/8003
```

The following example is sample output from the show bgp l2vpn vpls command with the process keyword:

```
RP/0/RP0/CPU0:router#show bgp l2vpn vpls process
Wed Mar 17 15:29:56.086 EDT

BGP Process Information:
BGP is operating in STANDALONE mode
Autonomous System number format: ASPLAIN
Autonomous System: 1
Router ID: 60.60.60.60 (manually configured)
Default Cluster ID: 60.60.60.60
Active Cluster IDs:  60.60.60.60
Fast external fallover enabled
Neighbor logging is enabled
Enforce first AS enabled
Default local preference: 100
Default keepalive: 60
Graceful restart enabled
Restart time: 180
Stale path timeout time: 1200
RIB purge timeout time: 900
Non-stop routing is enabled
Update delay: 600
Generic scan interval: 60

Address family: L2VPN VPLS
Dampening is not enabled
Client reflection is enabled in global config
Scan interval: 60
Main Table Version: 24001
Table version synced to RIB: 1
RIB has not converged: version 0
```

```
Node                 Process     Nbrs Estb Rst Upd-Rcvd Upd-Sent Nfn-Rcv Nfn-Snt
node0_RSP0_CPU0      Speaker        1    1   2    45347      237       0       0
```

show bgp neighbor-group

To display information about the Border Gateway Protocol (BGP) configuration for neighbor groups, use the **show bgp neighbor-group** command in XR EXEC mode.

show bgp neighbor-group *group-name* {**configuration** [**defaults**] [**nvgen**] | **inheritance** | **users**}

Syntax Description

group-name	Name of the address family group to display.
configuration	(Optional) Displays the effective configuration for the neighbor group, including any configuration inherited by this neighbor group.
defaults	(Optional) Displays all configuration, including default configuration.
nvgen	(Optional) Displays output in **show running-config** command output. If the **defaults** keyword is also specified, the output is not suitable for cutting and pasting into a configuration session.
inheritance	Displays the af-groups, session groups, and neighbor groups from which this neighbor group inherits configuration.
users	Displays the neighbors and neighbor groups that inherit configuration from this neighbor group.

Command Default

No default behavior or value

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

Use the **show bgp neighbor-group** command with the *group-name* **configuration** argument and keyword to display the effective configuration of a neighbor group, including any configuration inherited from session groups, address family groups, and neighbor groups through application of the **use** command. The source of each configured command is also displayed.

Use the **defaults** keyword to display all configuration for the neighbor group, including default configuration. The command output identifies default onfiguration. Use the **nvgen** keyword to display configuration in the output form of **show running-config**command. Output in this form is suitable for cutting and pasting into a configuration session.

The**show bgp neighbor-group** command with the *group-name* **inheritance** argument and keyword displays the session groups, address family groups, and neighbor groups from which the specified neighbor group inherits configuration.

The **show bgp neighbor-group** *group-name* command displays the neighbors and neighbor groups that inherit configuration from the specified neighbor group.

Task ID

Task ID	Operations
bgp	read

Examples

The examples use the following configuration:

```
 af-group group3 address-family ipv4 unicast
  remove-private-AS
  soft-reconfiguration inbound
!
 af-group group2 address-family ipv4 unicast
  use af-group group3
  send-community-ebgp
  send-extended-community-ebgp
  capability orf prefix both
 !
 session-group group3
  dmzlink-bw
 !
 neighbor-group group3
  use session-group group3
  timers 30 90
 !
 neighbor-group group1
  remote-as 1982
  use neighbor-group group2
  address-family ipv4 unicast
  !
 !
 neighbor-group group2
  use neighbor-group group3
  address-family ipv4 unicast
   use af-group group2
   weight 100
  !
```

The following is sample output from the **show bgp neighbor-group** command with the **configuration** keyword:

```
RP/0/RP0/CPU0:router# show bgp neighbor-group group1 configuration

neighbor-group group1
 remote-as 1982                   []
 timers 30 90                     [n:group2 n:group3]
 dmzlink-bw                       [n:group2 n:group3 s:group3]
 address-family ipv4 unicast      []
  capability orf prefix both      [n:group2 a:group2]
  remove-private-AS               [n:group2 a:group2 a:group3]
  send-community-ebgp             [n:group2 a:group2]
  send-extended-community-ebgp    [n:group2 a:group2]
  soft-reconfiguration inbound    [n:group2 a:group2 a:group3]
  weight 100                      [n:group2]
```

The configuration source is shown to the right of each command. In the output, the **remote-as** command is configured directly on neighbor group group1, and the **send-community-ebgp** command is inherited from neighbor group group2, which in turn inherits the setting from af-group group2.

The following is sample output from the **show bgp neighbor-group** command with the **users** keyword. This output shows that the group1 neighbor group inherits session (address family-independent configuration parameters) from the group2 neighbor group. The group1 neighbor group also inherits IPv4 unicast configuration parameters from the group2 neighbor group:

```
RP/0/RP0/CPU0:router# show bgp neighbor-group group2 users

Session:      n:group1
IPv4 Unicast: n:group1
```

The following is sample output from the **show bgp neighbor-group** command with the **inheritance** keyword. This output shows that the specified neighbor group group1 inherits session (address family-independent configuration) from neighbor group group2, which inherits its own session from neighbor group group3. Neighbor group group3 inherited its session from session group group3. It also shows that the group1 neighbor-group inherits IPv4 unicast configuration parameters from the group2 neighbor group, which in turn inherits them from the group2 af-group, which itself inherits them from the group3 af-group:

```
RP/0/RP0/CPU0:router# show bgp neighbor-group group1 inheritance

Session:      n:group2 n:group3 s:group3
IPv4 Unicast: n:group2 a:group2 a:group3
```

This table describes the significant fields shown in the display.

Table 16: show bgp neighbor-group Field Descriptions

Field	Description
[ ]	Configures the command directly on the specified address family group.
s:	Indicates the name that follows is a session group.
a:	Indicates the name that follows is an address family group.
n:	Indicates the name that follows is a neighbor group.
[dflt]	Indicates the setting is not explicitly configured or inherited, and the default value for the setting is used. This field may be shown when the **defaults** keyword is specified.
<not set>	Indicates that the default is for the setting to be disabled. This field may be shown when the **defaults** keyword is specified.

show bgp neighbors

To display information about Border Gateway Protocol (BGP) connections to neighbors, use the **show bgp neighbors** command in XR EXEC mode.

show bgp neighbors [**performance-statistics** | **missing-eor**] [**standby**] [**detail**]
show bgp neighbors *ip-address*[**advertised-routes** | **dampened-routes** | **flap-statistics** | **performance-statistics** | **received** | {**prefix-filter** | **routes**} | **routes**] [**standby**]
show bgp neighbors *ip-address* [**configuration** | [**defaults**] | **nvgen** | **inheritance**][**standby**]
show bgp neighbors *ip-address* [**decoded-message-log** | [**in**] | **out** | **standby**]

Syntax Description

ipv4	(Optional) Specifies IP Version 4 address prefixes.
unicast	(Optional) Specifies unicast address prefixes.
multicast	(Optional) Specifies multicast address prefixes.
labeled-unicast	(Optional) Specifies labeled unicast address prefixes.
all	(Optional) For subaddress families, specifies prefixes for all subaddress families.
tunnel	(Optional) Specifies tunnel address prefixes.
ipv6	(Optional) Specifies IP Version 6 address prefixes.
all	(Optional) For address family, specifies prefixes for all address families.
vpnv4 unicast	(Optional) Specifies VPNv4 unicast address families.
vrf	(Optional) Specifies VPN routing and forwarding (VRF) instance.
vrf-name	(Optional) Name of a VRF.
all	(Optional) For VRF, specifies all VRFs.
ipv4 { **unicast** | **labeled-unicast** }	(Optional) For VRF, specifies IPv4 unicast or labeled-unicast address families.
ipv6 unicast	(Optional) For VRF, specifies IPv6 unicast address families.
performance-statistics	(Optional) Displays performance statistics relative to work done by the BGP process for this neighbor.
missing-eor	(Optional) Displays neighbors that did not send end-of-rib (EoR) notification in read-only mode.
ip-address	(Optional) IP address of the BGP-speaking neighbor. If you omit this argument, all neighbors are displayed.
advertised-routes	(Optional) Displays all routes the router advertised to the neighbor.
dampened-routes	(Optional) Displays the dampened routes that are learned from the neighbor.

flap-statistics	(Optional) Displays flap statistics of the routes learned from the neighbor.
received { **prefix-filter** | **routes** }	(Optional) Displays information received from the BGP neighbor. The options are: **prefix-filter—** Displays the prefix list filter. **routes—**Displays routes from the neighbor before inbound policy
routes	(Optional) Displays routes learned from the neighbor.
configuration	(Optional) Displays the effective configuration for the neighbor, including any settings that have been inherited from session groups, neighbor groups, or af-groups used by this neighbor.
defaults	(Optional) Displays all configuration settings, including any default settings.
nvgen	(Optional) Displays output in the **show running-config** command output.
inheritance	(Optional) Displays the session groups, neighbor groups, and af-groups from which this neighbor inherits configuration settings.
in	(Optional) Displays BGP inbound messages.
out	(Optional) Displays BGP outbound messages.
standby	Displays standby BGP information.
detail	Displays detailed BGP information.

Command Default

If no address family or subaddress family is specified, the default address family and subaddress family specified using the **set default-afi** and **set default-safi** commands are used.

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 7.9.1	This command displays the slow peer configuration state and slow peer detection or processing information.
Release 6.0	This command was introduced.
Release 24.3.1	This command was modified. The show output was modified to display the BFD mode and status for the BGP session with the neighbor.

Usage Guidelines



Note The **set default-afi** command is used to specify the default address family for the session, and the **set default-safi** command is used to specify the default subaddress family for the session. See the for detailed information and syntax for the **set default-afi** and **set default-safi** commands. If you do not specify a default address family, the default address family is IPv4. If you do not specify a default subaddress family, the default subaddress family is unicast.

BGP contains a separate routing table for each configured address family and subaddress family combination. The address family and subaddress family options specify which routing table should be examined. If the **all** keyword is specified for address family or subaddress family, each matching routing table is examined in turn.

Use the **show bgp neighbors** command to display detailed information about all neighbors or a specific neighbor. Use the **performance-statistics** keyword to display information about the work related to specific neighbors done by the BGP process.

Use the **show bgp neighbors** command with the *ip-address* **received prefix-filter** argument and keyword to display the Outbound Route Filter (ORF) received from a neighbor.

Use the **advertised-routes** keyword to display a summary of the routes advertised to the specified neighbor.

Use the **dampened-routes** keyword to display routes received from the specified neighbor that have been suppressed due to dampening. For more details, see the **show bgp dampened-paths** command.

To display information about flapping routes received from a neighbor, use the **flap-statistics** keyword. For more details, see the **show bgp flap-statistics** command.

To display the routes received from a neighbor, use the **routes** keyword. For more details, see the **show bgp** command.

Use the **show bgp neighbor** command with the *ip-address* **configuration** argument and keyword to display the effective configuration of a neighbor, including configuration inherited from session groups, neighbor groups, or af-groups through application of the **use** command. Use the **defaults** keyword to display the value of all configurations for the neighbor, including default configuration. Use the **nvgen** keyword to display configuration output format of the **show running-config** command. Output in this format is suitable for cutting and pasting into a configuration session. Use the **show bgp neighbors** command with the *ip-address* **inheritance** argument and keyword to display the session groups, neighbor groups, and af-groups from which the specified neighbor inherits configuration.

Task ID

Task ID	Operations
bgp	read

The following is the sample output from the **show bgp neighbors** command with the *ip-address* and **configuration** argument and keyword to display graceful maintenance feature attributes:

```
**************************************************************************

Router# show bgp neighbor 10.12.12.5
...
  Graceful Maintenance locally active, Local Pref=45, AS prepends=3
...
```

```
 For Address Family: IPv4 Unicast
...
  GSHUT Community attribute sent to this neighbor
...

******************************************************************************

RP/0/0/CPU0:R1#show bgp neighbor 10.12.12.5 configuration
Mon Feb  2 14:30:41.042 PST
neighbor 10.12.12.5
 remote-as 1                          []
 graceful-maintenance 1               []
 gr-maint local-preference 45         []
 gr-maint as-prepends 3               []
 gr-maint activate                    []

******************************************************************************
```

Examples

The following is sample output from the **show bgp neighbors** command with BGP Persistence or long lived graceful restart (LLGR) status:

```
Router# show bgp neighbors 10.3.3.3
BGP neighbor is 10.3.3.3
Remote AS 30813, local AS 30813, internal link
Remote router ID 10.3.3.3
  BGP state = Established, up for 2d19h
  NSR State: NSR Ready
  BFD enabled (initializing)
  Last read 00:00:01, Last read before reset 2d19h
  Hold time is 180, keepalive interval is 60 seconds
  Configured hold time: 180, keepalive: 60, min acceptable hold time: 3
  Last write 00:00:03, attempted 19, written 19
  Second last write 00:01:03, attempted 19, written 19
  Last write before reset 2d19h, attempted 19, written 19
  Second last write before reset 2d19h, attempted 19, written 19
  Last write pulse rcvd  Nov 19 09:24:38.035 last full not set pulse count 66013
  Last write pulse rcvd before reset 2d19h
  Socket not armed for io, armed for read, armed for write
  Last write thread event before reset 2d19h, second last 2d19h
  Last KA expiry before reset 2d19h, second last 2d19h
  Last KA error before reset 00:00:00, KA not sent 00:00:00
  Last KA start before reset 2d19h, second last 2d19h
  Precedence: internet
  Non-stop routing is enabled
  Graceful restart is enabled
  Restart time is 120 seconds
  Stale path timeout time is 150 seconds
  Multi-protocol capability received
  Neighbor capabilities:
    Route refresh: advertised (old + new) and received (old + new)
    Graceful Restart (GR Awareness): advertised and received
    4-byte AS: advertised and received
    Address family IPv4 Unicast: advertised and received
    Address family VPNv4 Unicast: advertised and received
    Address family VPNv6 Unicast: advertised and received
    Address family RT Constraint: advertised and received
  Received 51634 messages, 0 notifications, 0 in queue
  Sent 33017 messages, 2 notifications, 0 in queue
  Minimum time between advertisement runs is 0 secs

For Address Family: IPv4 Unicast
```

```
  BGP neighbor version 204
  Update group: 0.2 Filter-group: 0.2  No Refresh request being processed
  AF-dependent capabilities:
    Graceful Restart capability advertised
      Local restart time is 120, RIB purge time is 600 seconds
      Maximum stalepath time is 150 seconds
    Graceful Restart capability received
      Remote Restart time is 120 seconds
      Neighbor preserved the forwarding state during latest restart
  Route refresh request: received 0, sent 0
  Policy for incoming advertisements is pass
  Policy for outgoing advertisements is pass
  0 accepted prefixes, 0 are bestpaths
  Cumulative no. of prefixes denied: 0.
  Prefix advertised 0, suppressed 0, withdrawn 0
  Maximum prefixes allowed 1048576
  Threshold for warning message 75%, restart interval 0 min
  AIGP is enabled
  An EoR was not received during read-only mode
  Last ack version 204, Last synced ack version 204
  Outstanding version objects: current 0, max 0
  Additional-paths operation: None
  Send Multicast Attributes

For Address Family: VPNv4 Unicast
  BGP neighbor version 8309
  Update group: 0.2 Filter-group: 0.2  No Refresh request being processed
  Inbound soft reconfiguration allowed
  AF-dependent capabilities:
    Graceful Restart capability advertised
      Local restart time is 120, RIB purge time is 600 seconds
      Maximum stalepath time is 150 seconds
    Graceful Restart capability received
      Remote Restart time is 120 seconds
      Neighbor preserved the forwarding state during latest restart
    Long-lived Graceful Restart Capability advertised
      Advertised Long-lived Stale time 3000 seconds
    Maximum acceptable long-lived stale time from this neighbor is 3000
    Long-lived Graceful Restart Capability received
      Received long-lived stale time is 3000 seconds
      Neighbor preserved the forwarding state during latest restart
  Route refresh request: received 0, sent 0
  Policy for incoming advertisements is pass
  Policy for outgoing advertisements is pass
  250 accepted prefixes, 250 are bestpaths
  Cumulative no. of prefixes denied: 0.
  Prefix advertised 100, suppressed 0, withdrawn 0
  Maximum prefixes allowed 2097152
  Threshold for warning message 75%, restart interval 0 min
  Peer will hold long-lived stale routes for 3000 seconds
  AIGP is enabled
  An EoR was not received during read-only mode
  Last ack version 8309, Last synced ack version 8309
  Outstanding version objects: current 0, max 1
  Additional-paths operation: None
  Send Multicast Attributes

For Address Family: VPNv6 Unicast
  BGP neighbor version 5
  Update group: 0.2 Filter-group: 0.2  No Refresh request being processed
  Inbound soft reconfiguration allowed
  AF-dependent capabilities:
    Graceful Restart capability advertised
      Local restart time is 120, RIB purge time is 600 seconds
```

```
      Maximum stalepath time is 150 seconds
    Graceful Restart capability received
      Remote Restart time is 120 seconds
      Neighbor preserved the forwarding state during latest restart
    Long-lived Graceful Restart Capability advertised
      Advertised Long-lived Stale time 3000 seconds
    Maximum acceptable long-lived stale time from this neighbor is 3000
    Long-lived Graceful Restart Capability received
      Received long-lived stale time is 3000 seconds
      Neighbor preserved the forwarding state during latest restart
  Route refresh request: received 0, sent 0
  Policy for incoming advertisements is pass
  Policy for outgoing advertisements is pass
  0 accepted prefixes, 0 are bestpaths
  Cumulative no. of prefixes denied: 0.
  Prefix advertised 0, suppressed 0, withdrawn 0
  Maximum prefixes allowed 1048576
  Threshold for warning message 75%, restart interval 0 min
  Peer will hold long-lived stale routes for 3000 seconds
  AIGP is enabled
  An EoR was not received during read-only mode
  Last ack version 5, Last synced ack version 5
  Outstanding version objects: current 0, max 0
  Additional-paths operation: None
  Send Multicast Attributes

For Address Family: RT Constraint
  BGP neighbor version 8
  Update group: 0.1 Filter-group: 0.1  No Refresh request being processed  RT constraint
nbr enabled for VPN updates:
  AF-dependent capabilities:
    Graceful Restart capability advertised
      Local restart time is 120, RIB purge time is 600 seconds
      Maximum stalepath time is 150 seconds
    Graceful Restart capability received
      Remote Restart time is 120 seconds
      Neighbor preserved the forwarding state during latest restart
    Long-lived Graceful Restart Capability advertised
      Advertised Long-lived Stale time 3000 seconds
    Maximum acceptable long-lived stale time from this neighbor is 3000
  Route refresh request: received 0, sent 0
  Policy for incoming advertisements is pass
  Policy for outgoing advertisements is pass
  1 accepted prefixes, 1 are bestpaths
  Cumulative no. of prefixes denied: 0.
  Prefix advertised 2, suppressed 0, withdrawn 0
  Maximum prefixes allowed 1048576
  Threshold for warning message 75%, restart interval 0 min
  Peer will hold long-lived stale routes for 3000 seconds
  AIGP is enabled
  An EoR was not received during read-only mode
  Last ack version 8, Last synced ack version 8
  Outstanding version objects: current 0, max 1
  Additional-paths operation: None
  Send Multicast Attributes

  Connections established 3; dropped 2
  Local host: 209.165.201.1, Local port: 179, IF Handle: 0x00000000
  Foreign host: 10.3.3.3, Foreign port: 62747
  Last reset 2d19h, due to BGP Notification sent: hold time expired
  Time since last notification sent to neighbor: 2d19h
  Error Code: hold time expired
  Notification data sent:
    None
```

The following is sample output from the **show bgp neighbors** command:

```
Router# show bgp neighbors 10.0.101.1

BGP neighbor is 10.0.101.1, remote AS 2, local AS 1, external link
 Description: routem neighbor
 Remote router ID 10.0.101.1
  BGP state = Established, up for 00:00:56
  TCP open mode: passive only
BGP neighbor is 198.51.100.253
 Remote AS 300, local AS 100, external link
 Remote router ID 0.0.0.0
  BGP state = Idle (LC/FIB for the neighbor in reloading)
  Last read 00:00:00, Last read before reset 00:05:12
  Hold time is 180, keepalive interval is 60 seconds
  Configured hold time: 180, keepalive: 60, min acceptable hold time: 3

  BFD enabled (session initializing)
  Last read 00:00:55, hold time is 180, keepalive interval is 60 seconds
  DMZ-link bandwidth is 1000 Mb/s
  Neighbor capabilities:
    Route refresh: advertised
    4-byte AS: advertised and received
    Address family IPv4 Unicast: advertised and received
    Address family IPv4 Multicast: advertised and received
  Received 119 messages, 0 notifications, 0 in queue
  Sent 119 messages, 22 notifications, 0 in queue
  Minimum time between advertisement runs is 60 seconds

 For Address Family: IPv4 Unicast
  BGP neighbor version 137
  Update group: 1.3
  Community attribute sent to this neighbor
  AF-dependant capabilities:
    Outbound Route Filter (ORF) type (128) Prefix-list:
      Send-mode: advertised
      Receive-mode: advertised
  Route refresh request: received 0, sent 0
  Policy for incoming advertisements is pass-all
  Policy for outgoing advertisements is pass-all
  5 accepted prefixes, 5 are bestpaths
  Prefix advertised 3, suppressed 0, withdrawn 0, maximum limit 1000000
  Threshold for warning message 75%


 For Address Family: IPv4 Multicast
  BGP neighbor version 23
  Update group: 1.2
  Route refresh request: received 0, sent 0
  Policy for incoming advertisements is pass-all
  Policy for outgoing advertisements is pass-all
  2 accepted prefixes, 2 are bestpaths
  Prefix advertised 0, suppressed 0, withdrawn 0, maximum limit 131072
  Threshold for warning message 75%

  Connections established 9; dropped 8
 Last reset 00:02:10, due to User clear requested (CEASE notification sent - administrative
reset)
  Time since last notification sent to neighbor: 00:02:10
  Error Code: administrative reset
  Notification data sent:
    None
```

This table describes the significant fields shown in the display.

Table 17: show bgp neighbors Field Descriptions

Field	Description
BGP neighbor	IP address of the BGP neighbor and its autonomous system number. If the neighbor is in the same autonomous system as the router, then the link between them is internal; otherwise, it is considered external.
Description	Neighbor specific description.
remote AS	• Number of the autonomous system to which the neighbor belongs. • Range for 2-byte Autonomous system numbers (ASNs) is 1 to 65535 . • Range for 4-byte Autonomous system numbers (ASNs) in asplain format is 1 to 4294967295. • Range for 4-byte Autonomous system numbers (ASNs) is asdot format is 1.0 to 65535.65535.
local AS	Autonomous system number of the local system. • Range for 2-byte Autonomous system numbers (ASNs) is 1 to 65535. • Range for 4-byte Autonomous system numbers (ASNs) in asplain format is 1 to 4294967295. • Range for 4-byte Autonomous system numbers (ASNs) is asdot format is 1.0 to 65535.65535.
internal link	Neighbor is an internal BGP peer.
external link	Neighbor is an external BGP peer.
Administratively shut down	Neighbor connection is disabled using the **shutdown** command.
remote router ID	Router ID (an IP address) of the neighbor.
Neighbor under common administration	Neighbor is internal or a confederation peer.
BGP state	Internal state of this BGP connection.
BFD enabled	Status of bidirectional forwarding detection.
TCP open mode	TCP mode used in establishing the BGP session. The following valid TCP mode are supported: • default—Accept active/passive connections • passive-only—Accept only passive connections • active-only—Accept only active connections initiated by the router

Field	Description
Last read	Time since BGP last read a message from this neighbor.
hold time	Hold time (in seconds) used on the connection with this neighbor.
keepalive interval	Interval for sending keepalives to this neighbor.
DMZ-link bandwidth	DMZ link bandwidth for this neighbor.
Neighbor capabilities	BGP capabilities advertised and received from this neighbor. The following valid BGP capabilities are supported: • Multi-protocol • Route refresh • Graceful restart • Outbound Route Filter (ORF) type (128) Prefix
Route refresh	Indicates that the neighbor supports dynamic soft reset using the route refresh capability.
4-byte AS	Indicates that the neighbor supports the 4-byte AS capability.
Address family	Indicates that the local system supports the displayed address family capability. If “received” is displayed, the neighbor also supports the displayed address family.
Received	Number of messages received from this neighbor, the number of notification messages received and processed from this neighbor, and the number of messages that have been received, but not yet processed.
Sent	Number of messages sent to this neighbor, the number of notification messages generated to be sent to this neighbor, and the number of messages queued to be sent to this neighbor.
Minimum time between advertisement runs	Advertisement interval (in seconds) for this neighbor.
For Address Family	Information that follows is specific to the displayed address family.
BGP neighbor version	Last version of the BGP database that was sent to the neighbor for the specified address family.
Update group	Update group to which the neighbor belongs.
Route reflector client	Indicates that the local system is acting as a route reflector for this neighbor.

Field	Description
Inbound soft reconfiguration allowed	Indicates that soft reconfiguration is enabled for routes received from this neighbor. **Note** If the neighbor has route refresh capability, then soft configuration received-only routes are not stored by the local system unless "override route refresh" is displayed.
eBGP neighbor with no inbound or outbound policy: defaults to drop	Indicates that the neighbor does not have an inbound or outbound policy configured using the **route-policy (BGP)** command. Hence, no routes are accepted from or advertised to this neighbor.
Private AS number removed from updates to this neighbor	Indicates that remove-private-AS is configured on the specified address family for this neighbor.
NEXT_HOP is always this router	Indicates that next-hop-self is configured on the specified address family for this neighbor.
Community attribute sent to this neighbor	Indicates that send-community-ebgp is configured on the specified address family for this neighbor.
Extended community attribute sent to this neighbor	Indicates that send-extended-community-ebgp is configured on the specified address family for this neighbor.
Default information originate	Indicates that default-originate is configured on the specified address family for this neighbor, together with the policy used, if one was specified in the default-originate configuration. An indication of whether the default route has been advertised to the neighbor is also shown.
AF-dependant capabilities	BGP capabilities that are specific to a particular address family. The following valid AF-dependent BGP capabilities are supported: • route refresh capability • route refresh capability OLD value
Outbound Route Filter	Neighbor has the Outbound Route Filter (ORF) capability for the specified address family. Details of the capabilities supported are also shown: Send-mode—"advertised" is shown if the local system can send an outbound route filter to the neighbor. "received" is shown if the neighbor can send an outbound route filter to the local system. Receive-mode—"advertised" is shown if the local system can receive an outbound route filter from the neighbor. "received" is shown if the neighbor can receive an outbound route filter from the local system.
Graceful Restart Capability	Indicates whether graceful restart capability has been advertised to and received from the neighbor for the specified address family.
Neighbor preserved the forwarding state during latest restart	Indicates that when the neighbor connection was last established, the neighbor indicated that it preserved its forwarding state for the specified address family.

Field	Description
Local restart time	Restart time (in seconds) advertised to this neighbor.
RIB purge time	RIB purge time (in seconds) used for graceful restarts.
Maximum stalepath time	Maximum time (in seconds) a path received from this neighbor may be marked as stale if the neighbor restarts.
Remote Restart time	Restart time received from this neighbor.
Route refresh request	Number of route refresh requests sent and received from this neighbor.
Outbound Route Filter (ORF)	"sent" indicates that an outbound route filter has been sent to this neighbor. "received" indicates that an outbound route filter has been received from this neighbor. **Note** A received outbound route filter may be displayed using the **show bgp neighbors** command with the **received prefix-filter** keywords.
First update is deferred until ORF or ROUTE-REFRESH is received	If the local system advertised the receive capability and the neighbor has advertised send capability, no updates are generated until specifically asked by the neighbor (using a ROUTE-REFRESH or ORF with immediate request).
Scheduled to send the Prefix-list filter	Indicates the local system is due to send an outbound route filter request in order to receive updates from the neighbor.
Inbound path policy	Indicates if an inbound path policy is configured.
Outbound path policy	Indicates if an outbound path policy is configured.
Incoming update prefix filter list	Indicates a prefix list is configured to filter inbound updates from the neighbor.
Default weight	Default weight for routes received from the neighbor.
Policy for incoming advertisements	Indicates a route policy is configured to be applied to inbound updates from the neighbor.
Policy for outgoing advertisements	Indicates a route policy is configured to be applied to outbound updates to the neighbor.
Type	Indicates whether the condition map selects routes that should be advertised, or routes that should not be advertised: Exist—Routes advertised if permitted by the condition route map. Non-exist—Routes advertised if denied by the condition route map.
accepted prefixes	Number of prefixes accepted.
Prefix advertised	Number of prefixes advertised to the neighbor during the lifetime of the current connection with the neighbor.

Field	Description
suppressed	Number of prefix updates that were suppressed because no transitive attributes changed from one best path to the next. **Note** Update suppression occurs only for external BGP neighbors.
withdrawn	Number of prefixes withdrawn from the neighbor during the lifetime of the current connection with the neighbor.
maximum limit	Maximum number of prefixes that may be received from the neighbor. If "(warning-only)" is displayed, a warning message is generated when the limit is exceeded, otherwise the neighbor connection is shut down when the limit is exceeded.
Threshold for warning message	Percentage of maximum prefix limit for the neighbor at which a warning message is generated.
Connections established	Number of times the router has established a BGP peering session with the neighbor.
dropped	Number of times that a good connection has failed or been taken down.
Last reset due to	Reason that the connection with the neighbor was last reset.
Time since last notification sent to neighbor	Amount of time since a notification message was last sent to the neighbor.
Error Code	Type of notification that was sent. The notification data, if any, is also displayed.
Time since last notification received from neighbor	Amount of time since a notification message was last received from the neighbor.
Error Code	Type of notification that was received. The notification data received, if any, is also displayed
External BGP neighbor may be up to <n> hops away	Indicates ebgp-multihop is configured for the neighbor.
External BGP neighbor not directly connected	Indicates that the neighbor is not directly attached to the local system.
Notification data sent:	Data providing more details on the error along with the error notification sent to the neighbor.

The following is sample output from the **show bgp neighbors** command with the **advertised-routes** keyword:

```
Router# show bgp neighbors 172.20.16.178 routes

BGP router identifier 172.20.16.181, local AS number 1
BGP main routing table version 27
BGP scan interval 60 secs
```

```
Status codes: s suppressed, d damped, h history, * valid, > best
              i - internal, S stale
Origin codes: i - IGP, e - EGP, ? - incomplete
Network            Next Hop            Metric LocPrf Weight Path
*> 10.0.0.0         172.20.16.178        40             0 10 ?
*> 10.22.0.0        172.20.16.178        40             0 10 ?
```

The following is sample output from the **show bgp neighbors** command with the **routes** keyword:

```
Router# show bgp neighbors 10.0.101.1 dampened-routes

BGP router identifier 10.0.0.5, local AS number 1
BGP main routing table version 48
Dampening enabled
BGP scan interval 60 secs
Status codes: s suppressed, d damped, h history, * valid, > best
              i - internal, S stale
Origin codes: i - IGP, e - EGP, ? - incomplete
   Network            From            Reuse    Path
*d 10.0.0.0           10.0.101.1      00:59:30 2 100 1000 i
*d 11.0.0.0           10.0.101.1      00:59:30 2 100 1000 i
*d 12.0.0.0           10.0.101.1      00:59:30 2 100 1000 i
*d 13.0.0.0           10.0.101.1      00:59:30 2 100 1000 i
*d 14.0.0.0           10.0.101.1      00:59:30 2 100 1000 i
```

This table describes the significant fields shown in the display.

Table 18: show bgp neighbors routes Field Descriptions

Field	Description
BGP router identifier	BGP identifier for the local system.
local AS number	Autonomous system number for the local system.
BGP main routing table version	Last version of the BGP database that was installed into the main routing table.
Dampening enabled	Displayed if dampening is enabled for the routes in this BGP routing table.
BGP scan interval	Interval (in seconds) between scans of the BGP table specified by the address family and subaddress family.

Field	Description
Status codes	Status of the table entry. The status is displayed as a three-character field at the beginning of each line in the table. The first character may be (in order of precedence): S—Path is stale, indicating that a graceful restart is in progress with the peer from which the route was learned. s—Path is more specific than a locally sourced aggregate route and has been suppressed. *—Path is valid. The second character may be (in order of precedence): >—Path is the best path to use for that network. d—Path is dampened. h—Path is a history entry, representing a route that is currently withdrawn, but that is being maintained to preserve dampening information. Such routes should never be marked as valid. The third character may be: i—Path was learned by an internal BGP (iBGP) session.
Origin codes	Origin of the path. The origin code is displayed at the end of each line in the table. It can be one of the following values: i—Path originated from an Interior Gateway Protocol (IGP) and was advertised with a **network** or **aggregate-address** command. e—Path originated from an Exterior Gateway Protocol (EGP). ?—Origin of the path is not clear. Usually, this is a route that is redistributed into BGP from an IGP.
Network	IP prefix and prefix length for a network.
Next Hop	IP address of the next system that is used when a packet is forwarded to the destination network. An entry of 0.0.0.0 indicates that the router has a non-BGP route to this network.
Metric	Value of the interautonomous system metric, otherwise known as the Multi Exit Discriminator (MED) metric.
LocPrf	Local preference value. This is used to determine the preferred exit point from the local autonomous system. It is propagated throughout the local autonomous system.
Weight	Path weight. Weight is used in choosing the preferred path to a route. It is not advertised to any neighbor.
Path	Autonomous system path to the destination network. At the end of the path is the origin code for the path.

The following is sample output from the **show bgp neighbors** command with the **dampened-routes** keyword:

```
Router# show bgp neighbors 10.0.101.1 flap-statistics

BGP router identifier 10.0.0.5, local AS number 1
BGP main routing table version 48
Dampening enabled
BGP scan interval 60 secs
Status codes: s suppressed, d damped, h history, * valid, > best
              i - internal, S stale
Origin codes: i - IGP, e - EGP, ? - incomplete
   Network            From            Flaps Duration Reuse    Path
 h 10.1.0.0           10.0.101.1      5008  2d02h             2 5000 1000
 h 10.2.0.0           10.0.101.1      5008  2d02h             2 2000 3000
 h 10.2.0.0           10.0.101.1      5008  2d02h             2 9000 6000
*d 10.0.0.0           10.0.101.1      5008  2d02h    00:59:30 2 100 1000
 h 10.0.0.0/16        10.0.101.1      5008  2d02h             2 100 102
*d 10.11.0.0          10.0.101.1      5008  2d02h    00:59:30 2 100 1000
*d 10.12.0.0          10.0.101.1      5008  2d02h    00:59:30 2 100 1000
*d 10.13.0.0          10.0.101.1      5008  2d02h    00:59:30 2 100 1000
*d 10.14.0.0          10.0.101.1      5008  2d02h    00:59:30 2 100 1000
 h 192.168.0.0/16     10.0.101.1      5008  2d02h             2 100 101
```

This table describes the significant fields shown in the display.

Table 19: show bgp neighbors dampened-routes Field Descriptions

Field	Description
BGP router identifier	BGP identifier for the local system.
local AS number	Autonomous system number for the local system.
BGP main routing table version	Last version of the BGP database that was installed into the main routing table.
Dampening enabled	Displayed if dampening is enabled for the routes in this BGP routing table.
BGP scan interval	Interval (in seconds) between scans of the BGP table specified by the address family and subaddress family.

Field	Description
Status codes	Status of the table entry. The status is displayed as a three-character field at the beginning of each line in the table. The first character may be (in order of precedence): S—Path is stale, indicating that a graceful restart is in progress with the peer from which the route was learned. s—Path is more specific than a locally sourced aggregate route and has been suppressed. *—Path is valid. The second character may be (in order of precedence): >—Path is the best path to use for that network. d—Path is dampened. h—Path is a history entry, representing a route that is currently withdrawn, but that is being maintained to preserve dampening information. Such routes should never be marked as valid. The third character may be: i—Path was learned by an internal BGP (iBGP) session.
Origin codes	Origin of the path. The origin code is displayed at the end of each line in the table. It can be one of the following values: i—Path originated from an Interior Gateway Protocol (IGP) and was advertised with a **network** or **aggregate-address** command. e—Path originated from an Exterior Gateway Protocol (EGP). ?—Origin of the path is not clear. Usually, this is a route that is redistributed into BGP from an IGP.
Network	IP prefix and prefix length for a network.
From	Neighbor from which the route was received.
Reuse	Time (in hours:minutes:seconds) after which the path is made available.
Path	Autonomous system path to the destination network. At the end of the path is the origin code for the path.

The following is sample output from the **show bgp neighbors** command with the **flap-statistics** keyword:

```
Router# show bgp neighbors 10.0.101.2 performance-statistics

BGP neighbor is 10.0.101.2, remote AS 1
  Read 3023 messages (58639 bytes) in 3019 calls (time spent: 1.312 secs)
  Read throttled 0 times
  Processed 3023 inbound messages (time spent: 0.198 secs)
  Wrote 58410 bytes in 6062 calls (time spent: 3.041 secs)
  Processing write list: wrote 0 messages in 0 calls (time spent: 0.000 secs)
  Processing write queue: wrote 3040 messages in 3040 calls (time spent: 0.055 secs)
```

```
Received 3023 messages, 0 notifications, 0 in queue
Sent 3040 messages, 0 notifications, 0 in queue
```

This table describes the significant fields shown in the display.

Table 20: show bgp neighbors flap-statistics Field Descriptions

Field	Description
BGP route identifier	BGP identifier for the local system.
local AS number	Autonomous system number for the local system.
BGP main routing table version	Last version of the BGP database that was installed into the main routing table.
Dampening enabled	Displayed if dampening has been enabled for the routes in this BGP routing table.
BGP scan interval	Interval (in seconds) between when the BGP process scans for the specified address family and subaddress family.
Status codes	Status of the table entry. The status is displayed as a three-character field at the beginning of each line in the table. The first character may be (in order of precedence): S—Path is stale, indicating that a graceful restart is in progress with the peer from which the route was learned. s—Path is more specific than a locally sourced aggregate route and has been suppressed. *—Path is valid. The second character may be (in order of precedence): d—Path is dampened. h—Path is a history entry, representing a route that is currently withdrawn, but that is being maintained to preserve dampening information. Such routes should never be marked as valid. The third character may be: i—Path was learned by an internal BGP (iBGP) session.
Origin codes	Origin of the path. The origin code is displayed at the end of each line in the table. It can be one of the following values: i—Path originated from an Interior Gateway Protocol (IGP) and was advertised with a **network** command. e—Path originated from an Exterior Gateway Protocol (EGP). ?—Origin of the path is not clear. Usually, this is a route that is redistributed into BGP from an IGP.
Network	IP prefix and prefix length for a network.

Field	Description
From	IP address of the peer that advertised this route.
Flaps	Number of times the route has flapped.
Duration	Time (in hours:minutes:seconds) since the router noticed the first flap.
Reuse	Time (in hours:minutes:seconds) after which the path is made available.
Path	Autonomous system path to reach the destination network.

The following is sample output from the **show bgp neighbors** command with the **performance-statistics** keyword:

```
Router# show bgp neighbors 10.0.101.2 performance-statistics
BGP neighbor is 10.0.101.2, remote AS 1
  Read 3023 messages (58639 bytes) in 3019 calls (time spent: 1.312 secs)
  Read throttled 0 times
  Processed 3023 inbound messages (time spent: 0.198 secs)
  Wrote 58410 bytes in 6062 calls (time spent: 3.041 secs)
  Processing write list: wrote 0 messages in 0 calls (time spent: 0.000 secs)
  Processing write queue: wrote 3040 messages in 3040 calls (time spent: 0.055 secs)
  Received 3023 messages, 0 notifications, 0 in queue
  Sent 3040 messages, 0 notifications, 0 in queue
```

This table describes the significant fields shown in the display.

Table 21: show bgp neighbors performance-statistics Field Descriptions

Field	Description
Read	Indicates the number of messages received from the neighbor, the total size of received messages, the number of read operations performed, and the real time spent (in seconds) by the process performing read operations for this neighbor.
Read throttled	Number of times that reading from the TCP connection to this neighbor has been throttled. Throttling is due to a backlog of messages that have been read but not processed.
inbound messages	Number of read messages that have been processed, and the real time spent processing inbound messages for this neighbor.
Wrote	Amount of data that has been sent to this neighbor, number of write operations performed, and the real time spent by the process performing write operations for this neighbor.
Processing write list	Number of messages written from the write list to this neighbor, number of times the write list has been processed, and real time spent processing the write list. **Note** Write lists typically contain only update messages.
Processing write queue	Number of messages written from the write queue to this neighbor, number of times the write queue has been processed, and real time spent processing the write queue.

Field	Description
Received	Number of messages received from this neighbor, number of notification messages received and processed from this neighbor, and number of messages that have been received, but not yet processed.
Sent	Number of messages sent to this neighbor, number of notification messages generated to be sent to this neighbor, and number of messages queued to be sent to this neighbor.

The following is sample output from the **show bgp neighbors** command with the **configuration** keyword:

```
Router# show bgp neighbors 10.0.101.1 configuration

neighbor 10.0.101.1
 remote-as 2                   []
 bfd fast-detect               []
 address-family ipv4 unicast   []
  policy pass-all in           []
  policy pass-all out          []
 address-family ipv4 multicast []
  policy pass-all in           []
  policy pass-all out          []
```

This table describes the significant fields shown in the display.

Table 22: show bgp neighbors configuration Field Descriptions

Field	Description
neighbor	IP address configuration of the neighbor.
remote-as	Remote autonomous system configured on the neighbor.
bfd fast-detect	BFD parameter configured on the neighbor.
address-family	Address family and subsequent address family configured on the router.
route-policy pass-all in	Route policy configured for inbound updates.
route-policy pass-all out	Route policy configured for outbound updates.

The following sample output shows sample output from **show bgp neighbors** command with additional paths send and receive capabilities advertised to neighbors:

```
BGP neighbor is 198.51.100.255
 Remote AS 100, local AS 100, internal link
 Remote router ID 33.33.33.33
  BGP state = Established, up for 19:54:12
  NSR State: None
  Last read 00:00:25, Last read before reset 19:54:54
  Hold time is 180, keepalive interval is 60 seconds
  Configured hold time: 180, keepalive: 60, min acceptable hold time: 3
  Last write 00:00:02, attempted 19, written 19
  Second last write 00:01:02, attempted 19, written 19
  Last write before reset 19:54:54, attempted 29, written 29
```

```
 Second last write before reset 19:54:59, attempted 19, written 19
 Last write pulse rcvd  Nov 11 12:58:03.838 last full not set pulse count 2407
 Last write pulse rcvd before reset 19:54:54
 Socket not armed for io, armed for read, armed for write
 Last write thread event before reset 19:54:54, second last 19:54:54
 Last KA expiry before reset 00:00:00, second last 00:00:00
 Last KA error before reset 00:00:00, KA not sent 00:00:00
 Last KA start before reset 19:54:54, second last 19:54:59
 Precedence: internet
 Non-stop routing is enabled
 Graceful restart is enabled
 Restart time is 120 seconds
 Stale path timeout time is 360 seconds
 Neighbor capabilities:            Adv         Rcvd
   Route refresh:                  Yes         Yes
   4-byte AS:                      Yes         Yes
   Address family IPv4 Unicast:    Yes         Yes
   Address family IPv4 Labeled-unicast:  Yes         Yes
   Address family VPNv4 Unicast:   Yes         Yes
   Address family IPv6 Unicast:    Yes         Yes
   Address family VPNv6 Unicast:   Yes         Yes
   Address family IPv4 MDT:        Yes         Yes
 Message stats:
   InQ depth: 0, OutQ depth: 0
                  Last_Sent               Sent  Last_Rcvd               Rcvd
   Open:          Nov 10 17:03:52.731        2  Nov 10 17:03:52.730        2
   Notification:  ---                        0  ---                        0
   Update:        Nov 10 17:05:02.435       20  Nov 10 17:04:58.812       12
   Keepalive:     Nov 11 12:58:03.632     1197  Nov 11 12:57:40.458     1196
   Route_Refresh: ---                        0  ---                        0
   Total:                                 1219                           1210
 Minimum time between advertisement runs is 0 secs

For Address Family: IPv4 Unicast
 BGP neighbor version 13
 Update group: 0.9
 NEXT_HOP is always this router
 AF-dependant capabilities:
   Graceful Restart capability advertised and received
     Neighbor preserved the forwarding state during latest restart
     Local restart time is 120, RIB purge time is 600 seconds
     Maximum stalepath time is 360 seconds
     Remote Restart time is 120 seconds
   Additional-paths Send: advertised and received
   Additional-paths Receive: advertised and received
 Route refresh request: received 0, sent 0
 0 accepted prefixes, 0 are bestpaths
 Prefix advertised 10, suppressed 0, withdrawn 0, maximum limit 524288
 Threshold for warning message 75%
 AIGP is enabled
 An EoR was received during read-only mode
 Last ack version 13, Last synced ack version 0
 Outstanding version objects: current 0, max 1
 Additional-paths operation: Send and Receive

For Address Family: IPv4 Labeled-unicast
 BGP neighbor version 13
 Update group: 0.4 (Update Generation Throttled)

 AF-dependant capabilities:
   Graceful Restart capability advertised and received
     Neighbor preserved the forwarding state during latest restart
     Local restart time is 120, RIB purge time is 600 seconds
     Maximum stalepath time is 360 seconds
```

```
    Remote Restart time is 120 seconds
  Additional-paths Send: received
  Additional-paths Receive: received
 Route refresh request: received 0, sent 0
 0 accepted prefixes, 0 are bestpaths
 Prefix advertised 2, suppressed 0, withdrawn 0, maximum limit 131072
 Threshold for warning message 75%
 AIGP is enabled
 An EoR was received during read-only mode
 Last ack version 13, Last synced ack version 0
 Outstanding version objects: current 0, max 1
 Additional-paths operation: None
```

This is sample output of the **show bgp neighbors** command when update wait-install is enabled. If the session open is postponed due to the reloading of the LC/FIB, the text "LC/FIB for the nieghobr in reloading" is displayed next to the BGP state.

```
Router#show bgp neighbors 198.51.100.253

BGP neighbor is 198.51.100.253
 Remote AS 300, local AS 100, external link
 Remote router ID 0.0.0.0
  BGP state = Idle (LC/FIB for the neighbor in reloading)
  Last read 00:00:00, Last read before reset 00:05:12
  Hold time is 180, keepalive interval is 60 seconds
  Configured hold time: 180, keepalive: 60, min acceptable hold time: 3
```

This is sample output from **show bgp neighbors** command that displays status of Accept Own configuration:

```
Router#show bgp neighbors 10.1.1.1

BGP neighbor is 10.1.1.1
 Remote AS 100, local AS 100, internal link
 Remote router ID 10.1.1.1
  BGP state = Established, up for 00:19:54
  NSR State: None
  Last read 00:00:55, Last read before reset 00:00:00
  Hold time is 180, keepalive interval is 60 seconds
  Configured hold time: 180, keepalive: 60, min acceptable hold time: 3
  Last write 00:00:54, attempted 19, written 19
  Second last write 00:01:54, attempted 19, written 19
  Last write before reset 00:00:00, attempted 0, written 0
  Second last write before reset 00:00:00, attempted 0, written 0
  Last write pulse rcvd  Jul 19 11:45:38.776 last full not set pulse count 43
  Last write pulse rcvd before reset 00:00:00
  Socket not armed for io, armed for read, armed for write
  Last write thread event before reset 00:00:00, second last 00:00:00
  Last KA expiry before reset 00:00:00, second last 00:00:00
  Last KA error before reset 00:00:00, KA not sent 00:00:00
  Last KA start before reset 00:00:00, second last 00:00:00
  Precedence: internet
  Non-stop routing is enabled
  Neighbor capabilities:
    Route refresh: advertised and received
    4-byte AS: advertised and received
    Address family VPNv4 Unicast: advertised and received
    Address family VPNv6 Unicast: advertised and received
  Received 22 messages, 0 notifications, 0 in queue
  Sent 22 messages, 0 notifications, 0 in queue
  Minimum time between advertisement runs is 0 secs
```

```
 For Address Family: VPNv4 Unicast

  BGP neighbor version 549
  Update group: 0.3 Filter-group: 0.1  No Refresh request being processed
  Route refresh request: received 0, sent 0
  Policy for incoming advertisements is pass-all
  Policy for outgoing advertisements is drop_111.x.x.x
  0 accepted prefixes, 0 are bestpaths
  Cumulative no. of prefixes denied: 0.
  Prefix advertised 0, suppressed 0, withdrawn 0
  Maximum prefixes allowed 524288
  Threshold for warning message 75%, restart interval 0 min
  AIGP is enabled
  Accept-own is enabled
  An EoR was received during read-only mode
  Last ack version 549, Last synced ack version 0
  Outstanding version objects: current 0, max 0
  Additional-paths operation: None

 For Address Family: VPNv6 Unicast

  BGP neighbor version 549
  Update group: 0.3 Filter-group: 0.1  No Refresh request being processed
  Route refresh request: received 0, sent 0
  Policy for incoming advertisements is pass-all
  Policy for outgoing advertisements is drop_111.x.x.x
  0 accepted prefixes, 0 are bestpaths
  Cumulative no. of prefixes denied: 0.
  Prefix advertised 0, suppressed 0, withdrawn 0
  Maximum prefixes allowed 524288
  Threshold for warning message 75%, restart interval 0 min
  AIGP is enabled
  Accept-own is enabled
  An EoR was received during read-only mode
  Last ack version 549, Last synced ack version 0
  Outstanding version objects: current 0, max 0
  Additional-paths operation: None

  Connections established 1; dropped 0
  Local host: 198.51.100.254, Local port: 179
  Foreign host: 10.1.1.1, Foreign port: 56391
  Last reset 00:00:00
RP/0/0/CPU0:BGP1-6#
```

This sample output from the **show bgp neighbor** command displays the status of permanent paths:

```
Router# show bgp neighbors 10.3.3.3
BGP neighbor is 10.3.3.3
Remote AS 30813, local AS 30813, internal link
Remote router ID 10.3.3.3
  BGP state = Established, up for 01:39:14
  Last read 00:00:58, Last read before reset 00:00:00
  Hold time is 180, keepalive interval is 60 seconds
  Configured hold time: 180, keepalive: 60, min acceptable hold time: 3
  Last write 00:00:53, attempted 2054, written 2054
  Second last write 00:00:53, attempted 45, written 45
  Last write before reset 00:00:00, attempted 0, written 0
  Second last write before reset 00:00:00, attempted 0, written 0
  Last write pulse rcvd  Aug 14 07:53:56.846 last full not set pulse count 226
  Last write pulse rcvd before reset 00:00:00
  Socket not armed for io, armed for read, armed for write
  Last write thread event before reset 00:00:00, second last 00:00:00
```

```
  Last KA expiry before reset 00:00:00, second last 00:00:00
  Last KA error before reset 00:00:00, KA not sent 00:00:00
  Last KA start before reset 00:00:00, second last 00:00:00
  Precedence: internet
  Multi-protocol capability received
  Neighbor capabilities:            Adv         Rcvd
    Route refresh:                  Yes         Yes
    4-byte AS:                      Yes         Yes
    Address family IPv4 Unicast:    Yes         Yes

For Address Family: IPv4 Unicast
  BGP neighbor version 1111
  Update group: 0.3 Filter-group: 0.5  No Refresh request being processed
  NEXT_HOP is always this router
  Default information originate: default sent
  AF-dependent capabilities:
    Additional-paths Send: received
    Additional-paths Receive: received
  Route refresh request: received 0, sent 0
  Policy for incoming advertisements is PASS
  Policy for outgoing advertisements is PASS
  100 accepted prefixes, 100 are bestpaths
  Cumulative no. of prefixes denied: 0.
  Prefix advertised 5500, suppressed 0, withdrawn 0
  Maximum prefixes allowed 1048576
  Threshold for warning message 75%, restart interval 0 min
  AIGP is enabled
  An EoR was received during read-only mode
  Last ack version 0, Last synced ack version 0
  Outstanding version objects: current 1, max 1
  Additional-paths operation: None
  Advertise Permanent-Network enabled

  Connections established 1; dropped 0
  Local host: 209.165.201.1, Local port: 179
  Foreign host: 10.3.3.3, Foreign port: 64742
  Last reset 00:00:00
```

The following is sample output from the **show bgp neighbors** command displaying BGP Monitoring Protocol (BMP) information:

```
Router# show bgp neighbors 10.1.1.2

Fri Sep 15 11:38:34.470 PST

BGP neighbor is 10.1.1.2
[...]
  Precedence: internet
  BGP Monitoring(BMP) activated for servers:
      2, 3
  Multi-protocol capability not received
[...]
```

The following is sample output from the **show bgp neighbors** command displaying BGP Persistence or long lived graceful restart (LLGR) status:

```
Router# show bgp neighbors 10.3.3.3

 For Address Family: VPNv4 Unicast
  BGP neighbor version 0
```

```
  Update group: 0.4 Filter-group: 0.0  No Refresh request being processed
  Inbound soft reconfiguration allowed
  Community attribute sent to this neighbor
  AF-dependent capabilities:
    Graceful Restart capability advertised
      Local restart time is 120, RIB purge time is 600 seconds
      Maximum stalepath time is 120 seconds
    Long-lived Graceful Restart Capability advertised
      Advertised Long-lived Stale time 16777215 seconds
    Maximum acceptable long-lived stale time from this neighbor is 16777215
    Treat neighbor as LLGR capable
    Remaining LLGR stalepath time 16776942
  Route refresh request: received 0, sent 0
```

This sample output from the **show bgp neighbor** command displays TCP MSS information for the specified neighbor:

```
Router# show bgp neighbor 10.0.0.2

BGP neighbor is 10.0.0.2
Remote AS 1, local AS 1, internal link
Remote router ID 10.0.0.2
BGP state = Established, up for 00:09:17
Last read 00:00:16, Last read before reset 00:00:00
Hold time is 180, keepalive interval is 60 seconds
Configured hold time: 180, keepalive: 60, min acceptable hold time: 3
Last write 00:00:16, attempted 19, written 19
Second last write 00:01:16, attempted 19, written 19
Last write before reset 00:00:00, attempted 0, written 0
Second last write before reset 00:00:00, attempted 0, written 0
Last write pulse rcvd Dec 7 11:58:42.411 last full not set pulse count 23
Last write pulse rcvd before reset 00:00:00
Socket not armed for io, armed for read, armed for write
Last write thread event before reset 00:00:00, second last 00:00:00
Last KA expiry before reset 00:00:00, second last 00:00:00
Last KA error before reset 00:00:00, KA not sent 00:00:00
Last KA start before reset 00:00:00, second last 00:00:00
Precedence: internet
Multi-protocol capability received
Neighbor capabilities:
Route refresh: advertised (old + new) and received (old + new)
Graceful Restart (GR Awareness): advertised and received
4-byte AS: advertised and received
Address family IPv4 Unicast: advertised and received
Received 12 messages, 0 notifications, 0 in queue
Sent 12 messages, 0 notifications, 0 in queue
Minimum time between advertisement runs is 0 secs
TCP Maximum Segment Size 500

For Address Family: IPv4 Unicast
BGP neighbor version 4
Update group: 0.2 Filter-group: 0.1 No Refresh request being processed
Route refresh request: received 0, sent 0
0 accepted prefixes, 0 are bestpaths
Cumulative no. of prefixes denied: 0.
Prefix advertised 0, suppressed 0, withdrawn 0
Maximum prefixes allowed 1048576
Threshold for warning message 75%, restart interval 0 min
AIGP is enabled
An EoR was received during read-only mode
Last ack version 4, Last synced ack version 0
```

```
Outstanding version objects: current 0, max 0
Additional-paths operation: None
Send Multicast Attributes
```

This sample output from the **show bgp neighbor** command with the **configuration** keyword displays TCP MSS configuration:

```
 Router# show bgp neighbor 10.0.0.2 configuration

neighbor 10.0.0.2
remote-as 1 []
tcp-mss 400 [n:n1]
address-family IPv4 Unicast []
```

The sample output from the **show bgp neighbor** with the **detail** keyword displays the neighbor address-family slow peer configuration state and slow peer detection or processing information.

```
Router# show bgp neighbors 198.51.100.254 detail
Thu Dec  1 02:40:40.301 UTC

BGP neighbor is 198.51.100.254
Remote AS 1, local AS 1, internal link
Remote router ID 198.51.100.254
Cluster ID 209.165.201.1
  BGP state = Established, up for 04:31:32
  Previous State: Active
  Last Received Message: Update
  NSR State: None
...
For Address Family: IPv6 Labeled-unicast
  BGP neighbor version 0
  Update group: 0.2 Filter-group: 0.60  No Refresh request being processed
  Route-Reflector Client
  Inbound soft reconfiguration allowed (override route-refresh)
  AF-dependent capabilities:
    Additional-paths Send: advertised
    Additional-paths Receive: advertised
  Route refresh request: received 0, sent 0
  Slow Peer State: Dynamic
    Detected state: TRUE, Detection threshold: 120
    Detection Count: 3, Recovery Count: 2
    Processing slow peer: FALSE
    Dynamic Trigger Count: 163, Dynamic Process Count: 1
```

Where:

- *Slow Peer State* indicates the effective configuration state (considering both global and neighbor address-family configuration) of neighbor address-family. *Slow Peer State* can be *Static* or *Dynamic* or *Detection-only* or *None*.
- *Detected State* indicates if the slow peer event is triggered for this neighbor address-family. *TRUE* if the neighbor address-family slow peer event is triggered, else *FALSE*.
- *Detection threshold* is the effective threshold configured (considering both global and neighbor address-family configuration) for the neighbor address-family. Threshold is used in determining if the peer (neighbor address-family) is slow.
- *Processing slow peer* indicates slow peer processing state of the neighbor address-family. *TRUE* if neighbor address-family is processed as slow peer else *FALSE*.

The bold text in sample output from the **show bgp neighbor** command with the **detail** keyword provides the information about the BFD mode and status for the BGP session with the neighbor at IP address 198.51.100.254. In the given example, BFD operates in the "BFD strict-mode-negotiate-override" mode. The other mode in which BFD may operate is: BFD strict-mode-negotiate.

```
Router# show bgp neighbors 198.51.100.254 detail
Thu Dec  1 02:40:40.301 UTC

BGP neighbor is 198.51.100.254
Wed Jun 12 10:53:10.957 UTC

BGP neighbor is 198.51.100.254
 Remote AS 1, local AS 1, internal link
 Remote router ID 3.3.3.3
 Cluster ID 2.2.2.2
  BGP state = Established, up for 00:01:35
  Previous State: Idle
  Last Received Message: KeepAlive
  NSR State: None
  BFD strict-mode-negotiate-override enabled (session up): mininterval: 333 multiplier: 3
  BFD discriminator: 0
  Last read 00:00:00, Last read before reset 00:00:00
  Hold time is 9, keepalive interval is 3 seconds
  Configured hold time: 9, keepalive: 3, min acceptable hold time: 3
  Last write 00:00:02, attempted 19, written 19
  Second last write 00:00:05, attempted 19, written 19
  Last write before reset 00:00:00, attempted 0, written 0
  Second last write before reset 00:00:00, attempted 0, written 0
  Last write pulse rcvd  Jun 12 10:53:10.763 last full not set pulse count 1424
  Last write pulse rcvd before reset 00:00:00
  Last insert into reset queue: not set, removed at not set
  Socket not armed for io, armed for read, armed for write
  Last write thread event before reset 00:00:00, second last 00:00:00
  Last KA expiry before reset 00:00:00, second last 00:00:00
  Last KA error before reset 00:00:00, KA not sent 00:00:00
  Last KA start before reset 00:00:00, second last 00:00:00
  First message received at Jun 12 10:51:35.359, sent at Jun 12 10:51:33.345
  Precedence: internet
  Non-stop routing is enabled
  Entered Neighbor NSR TCP mode:
    TCP Initial Sync :              ---
    TCP Initial Sync Phase Two :    ---
    TCP Initial Sync Done :         ---
  Scoped Sync Ilist Elem Flags : 0x00000000
  Multi-protocol capability received
  Neighbor capabilities:            Adv         Rcvd
    Route refresh:                  Yes         Yes
    4-byte AS:                      Yes         Yes
    Address family IPv4 Unicast:    Yes         Yes
    Address family VPNv4 Unicast:   Yes         Yes
    Address family VPNv6 Unicast:   Yes         Yes
    BFD Strict-Mode:                Yes         Yes
  Message stats:
    InQ depth: 0, OutQ depth: 0
                   Last_Sent               Sent  Last_Rcvd               Rcvd
    Open:          Jun 12 10:51:33.345        1  Jun 12 10:51:35.359        1
    Notification:  ---                        0  ---                        0
    Update:        ---                        0  Jun 12 10:52:22.969      485
    Keepalive:     Jun 12 10:53:08.362       32  Jun 12 10:53:10.763       19
    Route_Refresh: ---                        0  ---                        0
    Total:                                   33                           505
  Minimum time between advertisement runs is 0 secs
```

```
 Inbound message logging enabled, 3 messages buffered
 Outbound message logging enabled, 3 messages buffered
 Fast fallover is not enabled
   Neighbor is not directly connected
   Neighbor fast-fallover is not configured
 Event specific version bump counts:
   Interval definitions:
   Interval   Duration (min)  Start time            End time
   Live                    5  Jun 12 10:51:28.516   Jun 12 10:53:11.040
   1                       5  Jun 12 10:51:28.516   ---
   2                      15  Jun 12 10:51:28.516   ---
   3                      30  Jun 12 10:51:28.516   ---
   Category                            Total        Live  Interval 1  Interval 2  Interval
3
   Reachable                               0           0           0           0
  0
   Unreachable                             0           0           0           0
  0
```

show bgp nexthops

To display statistical information about the Border Gateway Protocol (BGP) next hops, use the **show bgp nexthops** command in XR EXEC mode.

show bgp nexthops A.B.C.D.aigp-value[**statistics**] [**speaker** *speaker-id*] [**standby**]

Syntax Description

ipv4	(Optional) Specifies IP Version 4 address prefixes.
unicast	(Optional) Specifies unicast address prefixes.
multicast	(Optional) Specifies multicast address prefixes.
labeled-unicast	(Optional) Specifies labeled-unicast address prefixes.
all	(Optional) For subaddress families, specifies prefixes for all subaddress families.
tunnel	(Optional) Specifies tunnel address prefixes.
ipv6	(Optional) Specifies IP Version 6 address prefixes.
all	(Optional) For address family, specifies prefixes for all address families.
vpnv4 unicast	(Optional) Specifies VPNv4 unicast address families.
vrf	(Optional) Specifies VPN routing and forwarding (VRF) instance.
vrf-name	(Optional) Name of a VRF.
all	(Optional) For VRF, specifies all VRFs.
ipv4 { **unicast** | **labeled-unicast** }	(Optional) For VRF, specifies IPv4 unicast or labeled-unicast address families.
ipv6 unicast	(Optional) For VRF, specifies IPv6 unicast address families.
statistics	(Optional) Specifies nexthop statistics.
speaker *speaker-id*	(Optional) Specifies a speaker process ID.
A.B.C.D	Next hop to display information about
aigp-value	Displays next hop statistics

Command Default

No default behavior or value

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

The **show bgp nexthops** command displays statistical information about next-hop notifications, the time spent processing the notifications, and details about each next-hop that has been registered with the Routing Information Base (RIB).

Use the **vrf** *vrf-name* keyword and argument to display only the next-hops present in the specified VPN routing and forwarding (VRF) instance.

The next-hop information is displayed for all active speaker processes in distributed mode. Each speaker displays a set of next-hops that belongs to the prefixes received by the speaker and next hops that belong to best paths that were received by other speaker processes. Use the **speaker** *speaker-id* keyword and argument to display information for only the specified speaker process. T he distributed mode must be defined using the **distributed speaker** command for the **speaker** keyword to be available.

Task ID

Task ID	Operations
bgp	read

Examples

The following is sample output from the **show bgp nexthops** command with the VRF specified:

```
RP/0/RP0/CPU0:router# show bgp vrf all nexthops

Fri Mar 13 17:05:40.656 UTC

VRF: 900
========

Total Nexthop Processing
  Time Spent: 0.000 secs

Maximum Nexthop Processing
  Received: 82y48w
  Bestpaths Deleted: 0
  Bestpaths Changed: 0
  Time Spent: 0.000 secs

Last Notification Processing
  Received: 1d22h
  Time Spent: 0.000 secs

IPv4 Unicast is active

Gateway Address Family: IPv4 Unicast
Table ID: 0xe0000001
Nexthop Count: 2
Critical Trigger Delay: 0msec
Non-critical Trigger Delay: 10000msec

Nexthop Version: 1, RIB version: 1

Status codes: R/UR Reachable/Unreachable
              C/NC Connected/Not-connected
              L/NL Local/Non-local
              I    Invalid (Policy Match Failed)
Next Hop        Status          Metric        Notf      LastRIBEvent RefCount
10.0.101.201    [UR]        4294967295         0/0       1d22h (Reg)        0/3
90.0.0.2        [R][C][NL]           0         1/0       1d22h (Cri)       20/23
```

```
VRF: 901
========

Total Nexthop Processing
  Time Spent: 0.000 secs

Maximum Nexthop Processing
  Received: 82y48w
  Bestpaths Deleted: 0
  Bestpaths Changed: 0
  Time Spent: 0.000 secs

Last Notification Processing
  Received: 1d22h
  Time Spent: 0.000 secs

IPv4 Unicast is active

Gateway Address Family: IPv4 Unicast
Table ID: 0xe0000002
Nexthop Count: 2
Critical Trigger Delay: 0msec
Non-critical Trigger Delay: 10000msec

Nexthop Version: 1, RIB version: 1

Status codes: R/UR Reachable/Unreachable
              C/NC Connected/Not-connected
              L/NL Local/Non-local
              I    Invalid (Policy Match Failed)
Next Hop        Status          Metric        Notf      LastRIBEvent RefCount
10.0.101.201    [UR]        4294967295         0/0       1d22h (Reg)        0/3
91.0.0.2        [R][C][NL]           0         1/0       1d22h (Cri)       10/13


VRF: 902
========

Total Nexthop Processing
  Time Spent: 0.000 secs

Maximum Nexthop Processing
  Received: 82y48w
  Bestpaths Deleted: 0
  Bestpaths Changed: 0
  Time Spent: 0.000 secs

Last Notification Processing
  Received: 1d22h
  Time Spent: 0.000 secs

IPv4 Unicast is active

Gateway Address Family: IPv4 Unicast
Table ID: 0xe0000003
Nexthop Count: 2
Critical Trigger Delay: 0msec
Non-critical Trigger Delay: 10000msec

Nexthop Version: 1, RIB version: 1

Status codes: R/UR Reachable/Unreachable
              C/NC Connected/Not-connected
              L/NL Local/Non-local
```

```
              I    Invalid (Policy Match Failed)
Next Hop        Status          Metric        Notf      LastRIBEvent RefCount
10.0.101.201    [UR]        4294967295         0/0       1d22h (Reg)        0/3
92.0.0.2        [R][C][NL]           0         1/0       1d22h (Cri)       10/13


VRF: 903
========

Total Nexthop Processing
  Time Spent: 0.000 secs

Maximum Nexthop Processing
  Received: 82y48w
  Bestpaths Deleted: 0
  Bestpaths Changed: 0
  Time Spent: 0.000 secs

Last Notification Processing
  Received: 1d22h
  Time Spent: 0.000 secs

IPv4 Unicast is active

Gateway Address Family: IPv4 Unicast
Table ID: 0xe0000004
Nexthop Count: 2
Critical Trigger Delay: 0msec
Non-critical Trigger Delay: 10000msec

Nexthop Version: 1, RIB version: 1

Status codes: R/UR Reachable/Unreachable
              C/NC Connected/Not-connected
              L/NL Local/Non-local
              I    Invalid (Policy Match Failed)
Next Hop        Status          Metric        Notf      LastRIBEvent RefCount
10.0.101.201    [UR]        4294967295         0/0       1d22h (Reg)        0/3
93.0.0.2        [R][C][NL]           0         1/0       1d22h (Cri)       10/13


VRF: 904
========

Total Nexthop Processing
  Time Spent: 0.000 secs

Maximum Nexthop Processing
  Received: 82y48w
  Bestpaths Deleted: 0
  Bestpaths Changed: 0
  Time Spent: 0.000 secs

Last Notification Processing
  Received: 1d22h
  Time Spent: 0.000 secs

IPv4 Unicast is active

Gateway Address Family: IPv4 Unicast
Table ID: 0xe0000005
Nexthop Count: 2
Critical Trigger Delay: 0msec
Non-critical Trigger Delay: 10000msec
```

```
Nexthop Version: 1, RIB version: 1

Status codes: R/UR Reachable/Unreachable
              C/NC Connected/Not-connected
              L/NL Local/Non-local
              I    Invalid (Policy Match Failed)
Next Hop        Status          Metric        Notf      LastRIBEvent RefCount
10.0.101.201    [UR]        4294967295         0/0       1d22h (Reg)        0/3
94.0.0.2        [R][C][NL]           0         1/0       1d22h (Cri)       10/13


VRF: 905
========

Total Nexthop Processing
  Time Spent: 0.000 secs

Maximum Nexthop Processing
  Received: 82y48w
  Bestpaths Deleted: 0
  Bestpaths Changed: 0
  Time Spent: 0.000 secs

Last Notification Processing
  Received: 1d22h
  Time Spent: 0.000 secs

IPv4 Unicast is active

Gateway Address Family: IPv4 Unicast
Table ID: 0xe0000006
Nexthop Count: 2
Critical Trigger Delay: 0msec
Non-critical Trigger Delay: 10000msec

Nexthop Version: 1, RIB version: 1

Status codes: R/UR Reachable/Unreachable
              C/NC Connected/Not-connected
              L/NL Local/Non-local
              I    Invalid (Policy Match Failed)
Next Hop        Status          Metric        Notf      LastRIBEvent RefCount
10.0.101.201    [UR]        4294967295         0/0       1d22h (Reg)        0/3
95.0.0.2        [R][C][NL]           0         1/0       1d22h (Cri)       10/13


VRF: 906
========

Total Nexthop Processing
  Time Spent: 0.000 secs

Maximum Nexthop Processing
  Received: 82y48w
  Bestpaths Deleted: 0
  Bestpaths Changed: 0
  Time Spent: 0.000 secs

Last Notification Processing
  Received: 1d22h
  Time Spent: 0.000 secs

IPv4 Unicast is active
```

```
Gateway Address Family: IPv4 Unicast
Table ID: 0xe0000007
Nexthop Count: 2
Critical Trigger Delay: 0msec
Non-critical Trigger Delay: 10000msec

Nexthop Version: 1, RIB version: 1

Status codes: R/UR Reachable/Unreachable
              C/NC Connected/Not-connected
              L/NL Local/Non-local
              I    Invalid (Policy Match Failed)
Next Hop        Status          Metric        Notf      LastRIBEvent RefCount
10.0.101.201    [UR]        4294967295         0/0       1d22h (Reg)        0/3
96.0.0.2        [R][C][NL]           0         1/0       1d22h (Cri)       10/13


VRF: 907
========

Total Nexthop Processing
  Time Spent: 0.000 secs

Maximum Nexthop Processing
  Received: 82y48w
  Bestpaths Deleted: 0
  Bestpaths Changed: 0
  Time Spent: 0.000 secs

Last Notification Processing
  Received: 1d22h
  Time Spent: 0.000 secs

IPv4 Unicast is active

Gateway Address Family: IPv4 Unicast
Table ID: 0xe0000008
Nexthop Count: 2
Critical Trigger Delay: 0msec
Non-critical Trigger Delay: 10000msec

Nexthop Version: 1, RIB version: 1

Status codes: R/UR Reachable/Unreachable
              C/NC Connected/Not-connected
              L/NL Local/Non-local
              I    Invalid (Policy Match Failed)
Next Hop        Status          Metric        Notf      LastRIBEvent RefCount
10.0.101.201    [UR]        4294967295         0/0       1d22h (Reg)        0/3
97.0.0.2        [R][C][NL]           0         1/0       1d22h (Cri)       10/13


VRF: 908
========

Total Nexthop Processing
  Time Spent: 0.000 secs

Maximum Nexthop Processing
  Received: 82y48w
  Bestpaths Deleted: 0
  Bestpaths Changed: 0
  Time Spent: 0.000 secs
```

```
Last Notification Processing
  Received: 1d22h
  Time Spent: 0.000 secs

IPv4 Unicast is active

Gateway Address Family: IPv4 Unicast
Table ID: 0xe0000009
Nexthop Count: 2
Critical Trigger Delay: 0msec
Non-critical Trigger Delay: 10000msec

Nexthop Version: 1, RIB version: 1

Status codes: R/UR Reachable/Unreachable
              C/NC Connected/Not-connected
              L/NL Local/Non-local
              I    Invalid (Policy Match Failed)
Next Hop        Status          Metric        Notf      LastRIBEvent RefCount
10.0.101.201    [UR]        4294967295         0/0       1d22h (Reg)        0/3
98.0.0.2        [R][C][NL]           0         1/0       1d22h (Cri)       10/13


VRF: 909
========

Total Nexthop Processing
  Time Spent: 0.000 secs

Maximum Nexthop Processing
  Received: 82y48w
  Bestpaths Deleted: 0
  Bestpaths Changed: 0
  Time Spent: 0.000 secs

Last Notification Processing
  Received: 1d22h
  Time Spent: 0.000 secs

IPv4 Unicast is active

Gateway Address Family: IPv4 Unicast
Table ID: 0xe000000a
Nexthop Count: 1
Critical Trigger Delay: 0msec
Non-critical Trigger Delay: 10000msec

Nexthop Version: 1, RIB version: 1

Status codes: R/UR Reachable/Unreachable
              C/NC Connected/Not-connected
              L/NL Local/Non-local
              I    Invalid (Policy Match Failed)
Next Hop        Status          Metric        Notf      LastRIBEvent RefCount
99.0.0.2        [UR]        4294967295         0/0       1d22h (Reg)        0/3


VRF: yellow
===========

Total Nexthop Processing
  Time Spent: 0.000 secs
```

```
Maximum Nexthop Processing
  Received: 82y48w
  Bestpaths Deleted: 0
  Bestpaths Changed: 0
  Time Spent: 0.000 secs

Last Notification Processing
  Received: 82y48w
  Time Spent: 0.000 secs

IPv4 Unicast is active

Gateway Address Family: IPv4 Unicast
Table ID: 0xe000000e
Nexthop Count: 0
Critical Trigger Delay: 0msec
Non-critical Trigger Delay: 10000msec

Nexthop Version: 1, RIB version: 1
```

This table describes the significant fields shown in the display.

Table 23: show bgp vrf all nexthops Field Descriptions

Field	Description
VRF	Name of the VRF.
Total Nexthop Processing Time Spent	Time spent processing trigger delays for critical and noncritical events for the VRF or address family. The time is specified in seconds.
Maximum Nexthop Processing	Time that has passed since the nexthop notification was received that resulted in spending the maximum amount of processing time for all notifications.
Last Notification Processing	Time that has passed since the last nexthop notification was received.
IPv4 Unicast is active.	VRF specified output that indicates the IPv4 unicast address family is active within the VRF.
Nexthop Count	Number of next hops for the VRF or address family.
Critical Trigger Delay	Configured critical trigger delay.
Non-critical Trigger Delay	Configured noncritical trigger delay.
Total Critical Notifications Received	Number of critical notifications received.
Total Non-critical Notifications Received	Number of noncritical notifications received.
Bestpaths Deleted After Last Walk	Number of best paths deleted due to the last notification.
Bestpaths Changed After Last Walk	Number of best paths modified due to the last notification.
Next Hop	IP address of the next hop.

Field	Description
Status	Status of the next hop.
Metric	IGP metric of the next hop.
Notf	Number of critical and noncritical notifications received.
LastRIBEvent	When the last notification was received from the RIB.
RefCount	The number of neighbors or prefixes that refer to the next hop in address family/all format.
Address Family	Name of the address family.

show bgp paths

To display all the Border Gateway Protocol (BGP) paths in the database, use the **show bgp paths** command in XR EXEC mode.

show bgp paths [**detail**] [**debug**] [**regexp** *regular-expression*]

Syntax Description

detail	(Optional) Displays detailed attribute information.
debug	(Optional) Displays attribute process ID, hash bucket, and hash chain ID attribute information.
regexp *regular-expression*	(Optional) Specifies an autonomous system path that matches the regular expression.

Command Default

No default behavior or values

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

Use the **show bgp paths** command to display information about AS paths and the associated attributes with which the paths were received.

If no options are specified, all stored AS paths are displayed with the number of routes using each path.

Note The AS path information is stored independently of the address family, making it possible that routes from different address families could be using the same path.

Use the *regular-expression* argument to limit the output to only those paths that match the specified regular expression. See the for information on regular expressions.

Use the **detail** keyword to display detailed information on the attributes stored with the AS path.

Task ID

Task ID	Operations
bgp	read

Examples

The following is sample output from the **show bgp paths** command:

```
RP/0/RP0/CPU0:router# show bgp paths detail

Proc  Attributes                    Refcount    Metric Path
Spk 0 ORG AS LOCAL                         7          0 i
```

```
Spk 0 ORG AS LOCAL COMM EXTCOMM            3          0 21 i
Spk 0 MET ORG AS                           3         55 2 i
Spk 0 ORG AS                               3          0 2 10 11 i
Spk 0 ORG AS COMM                          3          0 2 10 11 i
Spk 0 MET ORG AS ATOM                      3          2 2 3 4 ?
Spk 0 MET ORG AS                           3          1 2 3 4 e
Spk 0 MET ORG AS                           3          0 2 3 4 i
```

This table describes the significant fields shown in the display.

Table 24: show bgp paths Field Descriptions

Field	Description
Proc	ID of the process in which the path is stored. This is always “Spk 0.”
Attributes	Attributes that are present. The following may appear: MET —Multi Exit Discriminator (MED) attribute is present. ORG—Origin attribute is present. AS—AS path attribute is present. LOCAL—Local preference attribute is present. AGG—Aggregator attribute is present. COMM—Communities attribute is present. ATOM—Atomic aggregate attribute is present. EXTCOMM—Extended communities attribute is present.
NeighborAS	Autonomous system number of the neighbor, or 0, if the path information originated locally. • Range for 2-byte Autonomous system numbers (ASNs) is 1 to 65535. • Range for 4-byte Autonomous system numbers (ASNs) in asplain format is 1 to 4294967295. • Range for 4-byte Autonomous system numbers (ASNs) is asdot format is 1.0 to 65535.65535.
Refcount	Number of routes using a path.
Metric	Value of the interautonomous system metric, otherwise known as the MED metric.
Path	Autonomous system path to the destination network. At the end of the path is the origin code for the path: i—Path originated from an Interior Gateway Protocol (IGP) and was advertised with a **network** or **aggregate-address** command. e—Path originated from an Exterior Gateway Protocol (EGP). ?—Origin of the path is not clear. Usually, this is a route that is redistributed into BGP from an IGP.

show bgp policy

To display information about Border Gateway Protocol (BGP) advertisements under a proposed policy, use the **show bgp policy** command in XR EXEC mode.

show bgp policy

Syntax Description

ipv4	(Optional) Specifies IP Version 4 address prefixes.
unicast	(Optional) Specifies unicast address prefixes.
multicast	(Optional) Specifies multicast address prefixes.
labeled-unicast	(Optional) Specifies labeled unicast address prefixes.
all	(Optional) For subaddress families, specifies prefixes for all subaddress families.
tunnel	(Optional) Specifies tunnel address prefixes.
ipv6	(Optional) Specifies IP Version 6 address prefixes.
all	(Optional) For address family, specifies prefixes for all address families.
vpnv4 unicast	(Optional) Specifies VPNv4 unicast address families.
rd *rd-address*	(Optional) Displays routes with a specific route distinguisher.
vrf	(Optional) Specifies VPN routing and forwarding (VRF) instance.
vrf-name	(Optional) Name of a VRF.
all	(Optional) For VRF, specifies all VRFs.
ipv4 { **unicast** | **labeled-unicast** }	(Optional) For VRF, specifies IPv4 unicast or labeled-unicast address families.
ipv6 unicast	(Optional) For VRF, specifies IPv6 unicast address families.
neighbor	(Optional) Previews advertisements for a single neighbor.
ip-address	(Optional) IP address of a single neighbor.
sent-advertisements	(Optional) Displays the routes that have been advertised to neighbors. If a route has not yet been advertised to the neighbor, it is not shown.
route-policy	(Optional) Displays advertisements for an output route policy.
route-policy-name	(Optional) Name of the route policy.
summary	(Optional) Displays a summary of the BGP advertisements.

Command Default

Advertisements for all neighbors are displayed if the **neighbor** *ip-address* keyword and argument are not specified. If no address family or subaddress family is specified, the default address family and subaddress family specified using the **set default-afi** and **set default-safi** commands are used.

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines



Note The **set default-afi** command is used to specify the default address family for the session, and the **set default-safi** command is used to specify the default subaddress family for the session. See the for detailed information and syntax for the **set default-afi** and **set default-safi** commands. If you do not specify a default address family, the default address family is IPv4. If you do not specify a default subaddress family, the default subaddress family is unicast.

BGP contains a separate routing table for each configured address family and subaddress family combination. The address family and subaddress family options specify the routing table to be examined. If the **all** keyword is specified for the address family or subaddress family, each matching routing table is examined in turn.

Use the **show bgp policy** command to display routes that would be advertised to neighbors under a proposed policy. Unlike in the **show bgp advertised** command, the information displayed reflects any modifications made to the routes when executing the specified policy.

Use the **neighbor** keyword to limit the output to routes advertised to a particular neighbor. Use the **sent-advertisements** keyword to change the output in two ways:

- If a policy is not specified explicitly, any policy configured on the neighbor (using the **route-policy (BGP)** command) is executed before displaying the routes.
- Only routes that have already been advertised to the neighbor (and not withdrawn) are displayed. Routes that have not yet been advertised are not displayed.

Use the **summary** keyword to display abbreviated output.



Warning The **show bgp policy** command for any AFI/SAFI can increase CPU utilization under scale conditions. At high prefix scale, high neighbor scale, or both, the command can respond slowly and degrade concurrent BGP show operations.

Task ID

Task ID	Operations
bgp	read

Examples

The following is sample output from the **show bgp policy** command with the **summary** keyword in XR EXEC mode:

```
RP/0/RP0/CPU0:router# show bgp policy summary

Network           Next Hop        From            Advertised to
172.16.1.0/24     10.0.101.1      10.0.101.1      10.0.101.2
                                                  10.0.101.3

172.17.0.0/16     0.0.0.0         Local           10.0.101.1
                                                  10.0.101.2
                                                  10.0.101.3
```

This table describes the significant fields shown in the display.

Table 25: show bgp policy summary Field Descriptions

Field	Description
Network	IP prefix and prefix length for a network.
Next Hop	IP address of the next system that is used when a packet is forwarded to the destination network. An entry of 0.0.0.0 indicates that the router has a non-BGP route to this network.
From	IP address of the peer that advertised this route.
Local	Indicates the route originated on the local system.
Local Aggregate	Indicates the route is an aggregate created on the local system.
Advertised to	Indicates the neighbors to which this route was advertised.

The following is sample output from the **show bgp policy** command in XR EXEC mode:

```
RP/0/RP0/CPU0:router# show bgp policy

11.0.0.0/24 is advertised to 10.4.101.1
  Path info:
    neighbor: Local           neighbor router id: 10.4.0.1
    valid local best
  Attributes after inbound policy was applied:
    next hop: 0.0.0.0
    MET ORG AS
    origin: IGP metric: 0
    aspath:
  Attributes after outbound policy was applied:
    next hop: 10.4.0.1
    MET ORG AS
    origin: IGP metric: 0
    aspath: 1

11.0.0.0/24 is advertised to 10.4.101.2
  Path info:
    neighbor: Local           neighbor router id: 10.4.0.1
    valid local best
  Attributes after inbound policy was applied:
    next hop: 0.0.0.0
```

```
    MET ORG AS
    origin: IGP metric: 0
    aspath:
  Attributes after outbound policy was applied:
    next hop: 10.4.0.1
    MET ORG AS
    origin: IGP metric: 0
    aspath:

11.0.0.0/24 is advertised to 10.4.101.3
 Path info:
    neighbor: Local           neighbor router id: 10.4.0.1
    valid local best
  Attributes after inbound policy was applied:
    next hop: 0.0.0.0
    MET ORG AS
    origin: IGP metric: 0
    aspath:
  Attributes after outbound policy was applied:
    next hop: 10.4.0.1
    MET ORG AS
    origin: IGP metric: 0
    aspath:

12.0.0.0/24 is advertised to 10.4.101.2
 Path info:
    neighbor: 10.4.101.1       neighbor router id: 10.4.101.1
    valid external best
  Attributes after inbound policy was applied:
    next hop: 10.4.101.1
    ORG AS
    origin: IGP neighbor as: 2
    aspath: 2 3 4
  Attributes after outbound policy was applied:
    next hop: 10.4.101.1
    ORG AS
    origin: IGP neighbor as: 2
    aspath:2 3 4

12.0.0.0/24 is advertised to 10.4.101.3
 Path info:
    neighbor: 10.4.101.1       neighbor router id: 10.4.101.1
    valid external best
  Attributes after inbound policy was applied:
    next hop: 10.4.101.1
    ORG AS
    origin: IGP neighbor as: 2
    aspath: 2 3 4
  Attributes after outbound policy was applied:
    next hop: 10.4.101.1
    ORG AS
    origin: IGP neighbor as: 2
    aspath:2 3 4
```

This table describes the significant fields shown in the display.

Table 26: show bgp policy Field Descriptions

Field	Description
Is advertised to	IP address of the peer to which this route is advertised. If the route is advertised to multiple peers, information is shown separately for each peer.

Field	Description
neighbor	IP address of the peer that advertised this route, or one of the following: Local—Route originated on the local system. Local Aggregate—Route is an aggregate created on the local system.
neighbor router id	BGP identifier for the peer, or the local system if the route originated on the local system.
Not advertised to any peer	Indicates the no-advertise well-known community is associated with this route. Routes with this community are not advertised to any BGP peers.
Not advertised to any EBGP peer	Indicates the no-export well-known community is associated with this route. Routes with this community are not advertised to external BGP peers, even if those peers are in the same confederation as the local router.
Not advertised outside the local AS	Indicates the local-AS well-known community is associated with this route. Routes with this community value are not advertised outside the local autonomous system or confederation boundary.
(Received from a RR-client)	Path was received from a route reflector client.
(received-only)	Path is not used for routing purposes. It is used to support soft reconfiguration, and records the path attributes before inbound policy was applied to a path received from a peer. A path marked “received-only” indicates that either the path was dropped by inbound policy, or that a copy of path information was created and then modified for routing use.
(received & used)	Indicates that the path is used both for soft reconfiguration and routing purposes. A path marked “(received & used)”, implies the path information was not modified by inbound policy.
valid	Path is valid.
redistributed	Path is locally sourced through redistribution.
aggregated	Path is locally sourced through aggregation.
local	Path is locally sourced through the **network** command.
confed	Path was received from a confederation peer.
best	Path is selected as best.
multipath	Path is one of multiple paths selected for load-sharing purposes.

Field	Description
dampinfo	Indicates dampening information: Penalty—Current penalty for this path. Flapped—Number of times the route has flapped. In—Time (hours:minutes:seconds) since the network first flapped. Reuse in—Time (hours:minutes:seconds) after which the path is available. This field is displayed only if the path is currently suppressed.
Attributes after inbound policy was applied	Displays attributes associated with the received route, after any inbound policy has been applied. AGG—Aggregator attribute is present. AS—AS path attribute is present. ATOM—Atomic aggregate attribute is present. COMM—Communities attribute is present. EXTCOMM—Extended communities attribute is present. LOCAL—Local preference attribute is present. MET—Multi Exit Discriminator (MED) attribute is present. next hop—IP address of the next system used when a packet is forwarded to the destination network. An entry of 0.0.0.0 indicates that the router has a non-BGP route to this network. ORG—Origin attribute is present.
origin	Origin of the path: IGP—Path originated from an Interior Gateway Protocol (IGP) and was sourced by BGP using a **network** or **aggregate-address** command. EGP—Path originated from an Exterior Gateway Protocol. incomplete—Origin of the path is not clear; in example, a route that is redistributed into BGP from an IGP.
neighbor as	First autonomous system (AS) number in the AS path.
aggregator	Indicates that the path was received with the aggregator attribute. The AS number and router-id of the system that performed the aggregation are shown.
metric	Value of the interautonomous system metric, otherwise known as the MED metric.
localpref	Local preference value. This is used to determine the preferred exit point from the local autonomous system. It is propagated throughout the local autonomous system
aspath	AS path associated with the route.

Field	Description
community	Community attributes associated with the path. Community values are displayed in AA:NN format, except for the following well-known communities: Local-AS—Community with value 4294967043 or hex 0xFFFFFF03. Routes with this community value are not advertised outside the local autonomous system or confederation boundary. no-advertise—Community with value 4294967042 or hex 0xFFFFFF02. Routes with this community value are not advertised to any BGP peers. no-export—Community with value 4294967041 or hex 0xFFFFFF01. Routes with this community are not advertised to external BGP peers, even if those peers are in the same confederation as the local router.
Extended community	Extended community attributes associated with the path. For known extended community types, the following codes may be displayed: RT—Route target community SoO—Site of Origin community LB—Link Bandwidth community
Originator	Router ID of the originating router when route reflection is used.
Cluster lists	Router ID or cluster ID of all route reflectors through which the route has passed.
Attributes after outbound policy was applied	Displays attributes associated with the received route, after any outbound policy has been applied. AGG—Aggregator attribute is present. AS—AS path attribute is present. ATOM—Atomic aggregate attribute is present. COMM—Communities attribute is present. EXTCOMM—Extended communities attribute is present. LOCAL—Local preference attribute is present. MET—Multi Exit Discriminator (MED) attribute is present. next hop—IP address of the next system used when a packet is forwarded to the destination network. An entry of 0.0.0.0 indicates that the router has a non-BGP route to this network. ORG—Origin attribute is present.

show bgp process

To display Border Gateway Protocol (BGP) process information, use the **show bgp process** command in XR EXEC mode.

show bgp [**ipv4** | {**unicast** | **multicast** | **labeled-unicast** | **all** | **tunnel** | **mdt**} | **ipv6** | {**unicast** | **multicast** | **all** | **labeled-unicast**} | **all** | {**unicast** | **multicast** | **all** | **labeled-unicast** | **mdt** | **tunnel**} | **vpnv4 unicast** | **vpvn6 unicast**] **process** [**performance-statistics**] [**detail**] [**standby**]

Syntax Description

ipv4	(Optional) Specifies IP Version 4.
unicast	(Optional) Specifies the unicast subaddress family.
multicast	(Optional) Specifies the multicast subaddress family.
labeled-unicast	(Optional) Specifies labeled unicast address prefixes.
all	(Optional) For subaddress families, specifies prefixes for all subaddress families.
tunnel	(Optional) Specifies tunnel address prefixes.
ipv6	(Optional) Specifies IP Version 6.
all	(Optional) For address family, specifies prefixes for all address families.
vpnv4 unicast	(Optional) Specifies VPNv4 unicast address families.
performance- statistics	(Optional) Displays performance statistics relative to the work done by the specified process.
detail	(Optional) Specifies detailed process information.

Command Default

If no address family or subaddress family is specified, the default address family and subaddress family specified using the **set default-afi** and **set default-safi** commands are used.

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 7.5.3	Updated output to display the delay of the BGP start-up process since the last router update for Cisco 5500 series routers.
Release 6.0	This command was introduced.

Usage Guidelines



Note The **set default-afi** command is used to specify the default address family for the session, and the **set default-safi** command is used to specify the default subaddress family for the session. See the for detailed information and syntax for the **set default-afi** and **set default-safi** commands. If you do not specify a default address family, the default address family is IPv4. If you do not specify a default subaddress family, the default subaddress family is unicast.

Use the **show bgp process** command to display status and summary information for the Border Gateway Protocol (BGP) process. The output shows various global and address family-specific BGP configurations. A summary of the number of neighbors, update messages, and notification messages sent and received by the process is also displayed.

Use the **detail** keyword to display detailed process information. The detailed process information shows the memory used by each of various internal structure types.

Use the **performance-statistics** keyword to display a summary or detail of work done by the BGP processes. The summary display shows the real time spent performing certain operations and the time stamps for state transitions during initial convergence.

Task ID

Task ID	Operations
bgp	read

Examples

The following is sample output from the **show bgp process** command:

```
RP/0/RP0/CPU0:router# show bgp process

BGP Process Information
BGP is operating in STANDALONE mode
Autonomous System: 1
Router ID: 10.0.0.5 (manually configured)
Cluster ID: 10.0.0.5
Fast external fallover enabled
Neighbor logging is enabled
Enforce first AS enabled
Default local preference: 100
Default keepalive: 60
Update delay: 120
Generic scan interval: 60

Address family: IPv4 Unicast
Dampening is enabled
Client reflection is enabled
Scan interval: 60
Main Table Version: 150
IGP notification: IGPs notified

Node          Process     Nbrs Estab Rst Upd-Rcvd Upd-Sent Nfn-Rcvd Nfn-Sent
node0_0_CPU0 Speaker        3     2   1       20       10        0        0
```

This table describes the significant fields shown in the display.

Table 27: show bgp process Field Descriptions

Field	Description
BGP is operating in	Indicates BGP is operating in standalone mode. This is the only supported mode.
Autonomous System	Autonomous system number for the local system. • Range for 2-byte Autonomous system numbers (ASNs) is 1 to 65535. • Range for 4-byte Autonomous system numbers (ASNs) in asplain format is 1 to 4294967295. • Range for 4-byte Autonomous system numbers (ASNs) is asdot format is 1.0 to 65535.65535.
Router ID	BGP identifier assigned to the local system. If this is explicitly configured using the **bgp router-id** command, "manually configured" is displayed. If the router ID is not manually configured, it is determined from a global router ID. If no global ID is available, the router ID is shown as 0.0.0.0.
Confederation ID	Confederation identifier for the local system.
Cluster ID	Cluster identifier for the local system. If this is manually configured using the **bgp cluster-id** command, "manually configured" is displayed.
Default metric	Default metric. This is controlled by the **default-metric** command.
Fast external fallover enabled	Indicates whether fast external fallover is enabled. This is controlled by the **bgp fast-external-fallover disable** command.
Neighbor logging enabled	Indicates whether logging of peer connection up and down transitions is enabled. This is controlled by the **bgp log neighbor changes disable** command.
Enforce first AS enabled	Indicates that strict checking of the first AS number in paths received from external BGP peers is enabled. This is controlled by the **bgp enforce-first-as disable** command.
iBGP to IGP redistribution	Indicates internal redistribution is enabled using the **bgp redistribution-internal** command.
Treating missing MED as worst	Indicates missing Multi Exit Discriminator (MED) metric values are treated as worst in the route selection algorithm. This is controlled by the **bgp bestpath med missing-as-worst** command.
Always compare MED is enabled	Indicates that the MED is always used during the route selection algorithm, even when paths are received from external BGP neighbors in different autonomous systems. This is controlled by the **bgp bestpath med always** command.
AS Path ignore is enabled	Indicates that the AS path length is ignored by the route selection algorithm. This is controlled by the **bgp bestpath as-path ignore** command.

Field	Description
Comparing MED from confederation peers	Indicates that the MED values are used in the route selection algorithm when comparing routes received from confederation peers. This is controlled by the **bgp bestpath med confed** command.
Comparing router ID for eBGP paths	Indicates that the router ID is used as a tiebreaker by the route selection algorithm when comparing identical routes received from different external BGP neighbors. This is controlled by the **bgp bestpath compare-routerid** command.
Default local preference	Default local preference value used for BGP routes. This is controlled by the **bgp default local-preference** command.
Default keepalive	Default keepalive interval. This is controlled by the **timers bgp** command.
Graceful restart enabled	Indicates that the graceful restart capability is enabled. The configuration commands affecting graceful restart behavior are: • **bgp graceful-restart** • **bgp graceful-restart purge-time** • **bgp graceful-restart stalepath-time** • **bgp graceful-restart restart-time** • **bgp graceful-restart graceful-reset**
Update delay	Maximum time that a BGP process stays in read-only mode.
Generic scan interval	Interval (in seconds) between BGP scans for address family-independent tasks. This is controlled by the **bgp scan-time** command.
Dampening	Indicates whether dampening is enabled for the specified address family. This is controlled by the **dampening** command.
Client reflection	Indicates whether client-to-client route reflection is enabled for the specified address family. This is controlled by the **bgp client-to-client reflection disable** command.
Scan interval	Interval (in seconds) between BGP scans for the given address family. This is controlled by the **bgp scan-time** command in address family configuration mode.
Main Table Version	Last version of the BGP database that was installed into the main routing table.
IGP notification	Indicates whether Interior Gateway Protocols (IGP) have been notified of BGP convergence for the specified address family.
Node	Node on which the process is executing.
Process	Type of BGP process.
Speaker	Speaker process. A speaker process is responsible for receiving, processing, and sending BGP messages to configured neighbors.

Field	Description
Nbrs	Number of neighbors for which the process is responsible.
Estab	Number of neighbors that have connections in the established state for this process.
Rst	Number of times this process was restarted.
Upd-Rcvd	Number of update messages received by the process.
Upd-Sent	Number of update messages sent by the process.
Nfn-Rcvd	Number of notification messages received by the process.
Nfn-Sent	Number of notification messages sent by the process.

The following is sample output from the **show bgp process** command with the Graceful Maintenance feature enabled:

```
RP/0/0/CPU0:R1#show bgp process
...
Graceful Maintenance active. Retaining routes in RIB during BGP shutdown
...

Or

Graceful Maintenance active for all neighbors. Retaining routes in RIB during BGP shutdown

**********************************************************************

RP/0/0/CPU0:Jan 28 22:01:36.356 : bgp[1056]: %ROUTING-BGP-5-ADJCHANGE : neighbor 10.10.10.4
 Up (VRF: default) (AS: 4) WARNING: Graceful Maintenance is Active
```

The following is sample output from the **show bgp process** command with the **detail** keyword:

```
RP/0/RP0/CPU0:router# show bgp all all process detail

BGP Process Information
BGP is operating in STANDALONE mode
Autonomous System: 1
Router ID: 10.0.0.5 (manually configured)
Cluster ID: 10.0.0.5
Fast external fallover enabled
Neighbor logging is enabled
Enforce first AS enabled
Default local preference: 100
Default keepalive: 60
Update delay: 120
Generic scan interval: 60

BGP Speaker process: 0, location node0_0_0
Neighbors: 3, established: 2

                          Sent            Received
Updates:                  3               15
Notifications:            0               0

                          Number          Memory Used
Attributes:               12              1104
AS Paths:                 10              400
```

```
Communities:               2               1080
Extended communities:      1               40
Route Reflector Entries:   0               0
Route-map Cache Entries:   0               0
Filter-list Cache Entries: 0               0
Next Hop Cache Entries:    2               80
Update messages queued:    0

Address family: IPv4 Unicast
Dampening is enabled
Client reflection is enabled
Main Table Version: 12
IGP notification: IGPs notified

State: normal mode.
BGP Table Version: 12
Network Entries: 15, Soft Reconfig Entries: 0
Dampened Paths: 0, History Paths: 9

                           Allocated       Freed
Prefixes:                  15              0
Paths:                     19              0

                           Number          Memory Used
Prefixes:                  15              1230
Paths:                     19              760
```

This table describes the significant fields shown in the display.

Table 28: show bgp process detail Field Descriptions

Field	Description
BGP is operating in	Indicates whether BGP is operating in standalone mode.
Autonomous System	Autonomous system number for the local system.
Router ID	BGP identifier assigned to the local system. If this is explicitly configured using the **bgp router-id** command, “manually configured” is displayed. If the router ID is not manually configured, it is determined from a global router ID. If the global ID is not available, the router ID is shown as 0.0.0.0.
Confederation ID	Confederation identifier for the local system.
Cluster ID	Cluster identifier for the local system. If this is manually configured using the **bgp cluster-id** command, “manually configured” is displayed.
Default metric	Default metric.
Fast external fallover enabled	Indicates whether fast external fallover is enabled.
Neighbor logging enabled	Indicates whether logging of peer connection up and down transitions is enabled.
Enforce first AS enabled	Indicates that strict checking of the first autonomous system (AS) number in paths received from external BGP peers is enabled.

Field	Description
iBGP to IGP redistribution	Indicates internal redistribution is enabled using the **bgp redistribution-internal** command.
Treating missing MED as worst	Indicates missing MED metric values are treated as worst in the route selection algorithm. This is controlled by the **bgp bestpath med missing-as-worst** command.
Always compare MED is enabled	Indicates that the MED is always used during the route selection algorithm, even when paths are received from external BGP neighbors in different autonomous systems. This is controlled by the **bgp bestpath med always** command.
AS Path ignore is enabled	Indicates that the AS path length is ignored by the route selection algorithm. This is controlled by the **bgp bestpath as-path ignore** command.
Comparing MED from confederation peers	Indicates that the MED values are used in the route selection algorithm when comparing routes received from confederation peers. This is controlled by the **bgp bestpath med confed** command.
Comparing router ID for eBGP paths	Indicates that the router ID is used as a tiebreaker by the route selection algorithm when comparing identical routes received from different external BGP neighbors. This is controlled by the **bgp bestpath compare-routerid** command.
Default local preference	Default local preference value used for BGP routes.
Default keepalive	Default keepalive interval. This is controlled by the **timers bgp** command.
Graceful restart enabled	Indicates that the graceful restart capability is enabled. The configuration commands affecting graceful restart behavior are: • **bgp graceful-restart** • **bgp graceful-restart purge-time** • **bgp graceful-restart stalepath-time** • **bgp graceful-restart restart-time** • **bgp graceful-restart graceful-reset**
Update delay	Maximum time that a BGP process stays in read-only mode.
Generic scan interval	Interval (in seconds) between BGP scans for address family-independent tasks. This is controlled by the **bgp scan-time** command.
BGP Speaker Process	Speaker process responsible for receiving, processing and sending BGP messages.
Node	Node on which the specified process is executing.
Neighbors	Number of neighbors for which the specified process is responsible.
established	Number of neighbors that have connections in the established state for the specified process.

Field	Description
Updates	Number of update messages sent and received by the specified process.
Notifications	Number of notification messages sent and received by the specified process.
Attributes	Number of unique sets of attribute information stored in the specified process and the amount of memory used by the attribute information.
AS Paths	Number of unique autonomous system paths stored in the specified process and the amount of memory used by the AS path information.
Communities	Number of unique sets of community information stored in the specified process and the amount of memory used by them.
Extended communities	Number of unique sets of extended community information stored in the specified process and the amount of memory used by them.
Route Reflector Entries	Number of unique sets of route reflector information stored in the specified process and the amount of memory used by them.
Nexthop Entries	Number of entries and memory usage for cached next- hop information.
Update messages queued	Total number of update messages queued to be sent across all neighbors for which the specified process is responsible.
Address family	Specified address family.
Dampening	Indicates whether dampening is enabled for the specified address family.
Client reflection	Indicates whether client-to-client route reflection is enabled for the specified address family. This is controlled by the **bgp client-to-client reflection disable** command.
Scan interval	Interval (in seconds) between BGP scans for the given address family. This is controlled by the **bgp scan-time** command.
Main Table Version	Last version of the local BGP database for the specified address family that was injected into the main routing table.
IGP notification	Indicates whether IGPs have been notified of BGP convergence for the specified address family.
RIB has converged	Indicates whether the main routing table version has converged and the version at which it converged.

Field	Description
State	BGP system state for the specified address family and process. This may be one of the following: read-only mode—Initial set of updates is being recovered. In this mode, route selection is not performed, routes are not installed in the global RIB, and updates are not advertised to peers. best-path calculation mode—Route selection is being performed for the routes that were received while in read-only mode. import mode—Routes are imported from one VRF to another VRF once the best paths are calculated. This mode is supported in VPNv4 unicast address family mode. RIB update mode—Routes that were selected in best-path calculation mode are being installed in the global RIB. label allocation mode: Labels are allocated for the received prefixes based on the requirement. normal mode—Best paths are sent to the peers for routes that exist in the RIB. The route selection, import processing, RIB updates, and label allocation are performed as new updates are received.
BGP Table Version	Last version used in the BGP database for received routes.
Attribute download	Indicates whether the RIB attribute download is enabled.
Network Entries	Number of sets of prefix information held in the specified BGP process for the specified address family.
Soft Reconfig Entries	Number of sets of prefix information that are present only for the purpose of supporting soft reconfiguration.
Dampened Paths	Number of routes that are suppressed due to dampening for the specified address family.
History Paths	Number of routes that are currently withdrawn, but are being maintained to preserve dampening information.
Prefixes (Allocated/Freed)	Number of sets of prefix information for the specified address family that have been allocated and freed during the lifetime of the process.
Paths (Allocated/Freed)	Number of sets of route information for the specified address family that have been allocated and freed during the lifetime of the process.
Prefixes (Number/Memory Used)	Number of sets of prefix information currently allocated for the specified address family, and the amount of memory used by them.
Paths (Number/Memory Used)	Number of sets of route information currently allocated for the specified address family, and the amount of memory used by them.

The following is sample output from the **show bgp process** command with the **performance-statistics** keyword:

```
RP/0/RP0/CPU0:router# show bgp process performance-statistics detail

BGP Speaker process: 0, Node: node0_0_CPU0
Restart count: 2
Neighbors: 3, established: 2

                           Sent            Received
Updates:                   20              20
Notifications:             0               0

                           Number          Memory Used
Attributes:                2               184
AS Paths:                  2               48
Communities:               0               0
Extended communities:      0               0
Route Reflector Entries:   0               0
Route-map Cache Entries:   0               0
Filter-list Cache Entries: 0               0
Next Hop Cache Entries:    2               80
Update messages queued:    0

Read 14 messages (1142 bytes) in 12 calls (time spent: 0.024 secs)
Read throttled 0 times
Processed 14 inbound messages (time spent: 0.132 secs)
Wrote 2186 bytes in 24 calls (time spent: 0.024 secs)
Processing write list: wrote 18 messages in 4 calls (time spent: 0.000 secs)
Processing write queue: wrote 10 messages in 20 calls (time spent: 0.000 secs)
Socket setup (LPTS): 4 calls (time spent: 0.010 secs)
Configuration: 1 requests (time spent: 0.002 secs)
Operational data: 9 requests (time spent: 0.026 secs)


State: normal mode.
BGP Table Version: 150
Network Entries: 149, Soft Reconfig Entries: 0



                           Allocated       Freed
Prefixes:                  149             0
Paths:                     200             0

                           Number          Memory Used
Prefixes:                  149             12516
Paths:                     200             8000

Updates generated: 149 prefixes in 8 messages from 2 calls (time spent: 0.046 secs)
Scanner: 2 scanner runs (time spent: 0.008 secs)
RIB update: 1 rib update runs, 149 prefixes installed (time spent: 0.024 secs)
Process has converged for IPv4 Unicast.

 Update wait-install enabled:
          ack request 2, ack rcvd 2, slow ack 0
        Max time for batched RIB update:
          update time 0 secs, 1752000 nsecs
          install time 0 secs, 36391000 nsecs
          routes installed 6, modified 0, withdrawn 2
          start version 1, target version 15


First neighbor established: 1082604050s
Entered DO_BESTPATH mode: 1082604055s
Entered DO_RIBUPD mode: 1082604055s
```

```
Entered Normal mode: 1082604055s
Latest UPDATE sent: 1082604056s
```

This table describes the significant fields shown in the display.

Table 29: show bgp process performance-statistics Field Descriptions

Field	Description
BGP is operating in	Indicates whether BGP is operating in standalone mode.
Autonomous system	Autonomous system number for the local system.
Router ID	BGP identifier assigned to the local system. If this is explicitly configured using the **bgp router-id** command, “manually configured” is displayed. If the router ID is not manually configured, it is determined from a global router ID. If the global ID is not available, the router ID is shown as 0.0.0.0.
Confederation ID	Confederation identifier for the local system.
Cluster ID	The cluster identifier for the local system. If this is manually configured using the **bgp cluster-id** command, “manually configured” is displayed.
Default metric	Default metric.
Fast external fallover enabled	Indicates whether fast external fallover is enabled.
Neighbor logging enabled	Indicates whether logging of peer connection up and down transitions is enabled. This is controlled by the **bgp log neighbor changes disable** command.
Enforce first AS enabled	Indicates that strict checking of the first AS number in paths received from external BGP peers is enabled.
iBGP to IGP redistribution	Indicates internal redistribution is enabled using the **bgp redistribution-internal** command.
Treating missing MED as worst	Indicates missing MED metric values are treated as worst in the route selection algorithm. This is controlled using the **bgp bestpath med missing-as-worst** command.
Always compare MED is enabled	Indicates that the MED is always used during the route selection algorithm, even when paths are received from external BGP neighbors in different autonomous systems. This setting is controlled by the **bgp bestpath med always** command.
AS Path ignore is enabled	Indicates that the AS path length is ignored by the route selection algorithm. This is controlled by the **bgp bestpath as-path ignore** command.
Comparing MED from confederation peers	Indicates that the MED values are used in the route selection algorithm when comparing routes received from confederation peers. This is controlled by the **bgp bestpath med confed** command.

Field	Description
Comparing router ID for eBGP paths	Indicates that the router ID is used as a tiebreaker by the route selection algorithm when comparing identical routes received from different external BGP neighbors. This is controlled by the **bgp bestpath compare-routerid** command.
Default local preference	Default local preference value used for BGP routes.
Default keepalive	Default keepalive interval. This setting is controlled by the **timers bgp** command.
Graceful restart enabled	Indicates that the graceful restart capability is enabled. The configuration commands affecting graceful restart behavior are: **bgp graceful-restart**, **bgp graceful-restart purge-time**, **bgp graceful-restart stalepath-time**, **bgp graceful-restart restart-time**, and **bgp graceful-restart graceful-reset**.
Update delay	Maximum time that a BGP process stays in read-only mode.
Generic scan interval	Interval (in seconds) between BGP scans for address family-independent tasks. This setting is controlled by the **bgp scan-time** command in router configuration mode.
Address family	Specified address family.
Dampening	Indicates whether dampening is enabled for the specified address family.
Client reflection	Indicates whether client-to-client route reflection is enabled for the specified address family. This is controlled by the **bgp client-to-client reflection disable** command.
Scan interval	Interval (in seconds) between BGP scans for the given address family. This is controlled by the **bgp scan-time** command.
Main Table Version	Last version of the local BGP database for the specified address family that was injected into the main routing table.
IGP notification	Indicates whether IGPs have been notified of BGP convergence for the specified address family.
Node	Node on which the process is executing.
Process	BGP process.
Speaker	Speaker process. The speaker process is responsible for receiving, processing and sending BGP messages.
Read	Real time (in seconds) spent reading messages from peers by this process.
Write	Real time (in seconds) spent writing messages to peers by this process.
Inbound	The real time (in seconds) spent processing messages read from peers by this process.

Field	Description
Config	Real time (in seconds) spent processing configuration commands by this process.
Data	Real time (in seconds) spent providing operational data by this process.
Conv	Indicates whether the process has converged after the initial update.
Nbr Estab	Time stamp (in seconds) recording the time when the first neighbor became established.
Bestpath	Time stamp (in seconds) recording the time the best-path calculation mode was entered.
RIB Inst	Time stamp (in seconds) recording the time RIB update mode was entered.
Read/Write	Time stamp (in seconds) recording the time normal mode was entered.
Last Upd	Time stamp (in seconds) recording the time the last update was sent to a neighbor.
Address Family IPv4 Unicast converged in *n* seconds	Indicates that BGP has reached initial convergence for the IPv4 unicast address family. The time taken for convergence is shown.
Address Family IPv6 Multicast converged in *n* seconds	Indicates that BGP has reached initial convergence for the IPv6 multicast address family. The time taken for convergence is shown.
Update wait-install enabled	Indicates the **update wait-install** was configured.

The following is sample output from the **show bgp process** command with the **performance-statistics** and **detail** keywords:

```
RP/0/RP0/CPU0:router# show bgp process performance-statistics detail

BGP Speaker process: 0, Node: node0_0_CPU0
Restart count: 2
Neighbors: 3, established: 2

                           Sent            Received
Updates:                   20              20
Notifications:             0               0

                           Number          Memory Used
Attributes:                2               184
AS Paths:                  2               48
Communities:               0               0
Extended communities:      0               0
Route Reflector Entries:   0               0
Route-map Cache Entries:   0               0
Filter-list Cache Entries: 0               0
Next Hop Cache Entries:    2               80
Update messages queued:    0

Read 14 messages (1142 bytes) in 12 calls (time spent: 0.024 secs)
Read throttled 0 times
Processed 14 inbound messages (time spent: 0.132 secs)
Wrote 2186 bytes in 24 calls (time spent: 0.024 secs)
```

```
Processing write list: wrote 18 messages in 4 calls (time spent: 0.000 secs)
Processing write queue: wrote 10 messages in 20 calls (time spent: 0.000 secs)
Socket setup (LPTS): 4 calls (time spent: 0.010 secs)
Configuration: 1 requests (time spent: 0.002 secs)
Operational data: 9 requests (time spent: 0.026 secs)


State: normal mode.
BGP Table Version: 150
Network Entries: 149, Soft Reconfig Entries: 0



                          Allocated       Freed
Prefixes:                 149             0
Paths:                    200             0

                          Number          Memory Used
Prefixes:                 149             12516
Paths:                    200             8000

Updates generated: 149 prefixes in 8 messages from 2 calls (time spent: 0.046 secs)
Scanner: 2 scanner runs (time spent: 0.008 secs)
RIB update: 1 rib update runs, 149 prefixes installed (time spent: 0.024 secs)
Process has converged for IPv4 Unicast.

First neighbor established: 1082604050s
Entered DO_BESTPATH mode: 1082604055s
Entered DO_RIBUPD mode: 1082604055s
Entered Normal mode: 1082604055s
Latest UPDATE sent: 1082604056s
```

This table describes the significant fields shown in the display.

Table 30: show bgp process performance-statistics detail Field Descriptions

Field	Description
Process	The specified process.
Location	Node in which the specified process is executing.
Neighbors	Number of neighbors for which the specified process is responsible.
established	Number of neighbors that have connections in the established state for the specified process.
Updates	Number of update messages sent and received by the specified process.
Notifications	Number of notification messages sent and received by the specified process.
Attributes	Number of unique sets of attribute information stored in the specified process and the amount of memory used by the attribute information.
AS Paths	Number of unique autonomous system paths stored in the specified process, and the amount of memory used by the AS path information.
Communities	Number of unique sets of community information stored in the specified process and the amount of memory used by them.

Field	Description
Extended communities	Number of unique sets of extended community information stored in the specified process and the amount of memory used by them.
Route Reflector Entries	Number of unique sets of route reflector information stored in the specified process and the amount of memory used by them.
Route-map Cache Entries	Number of entries and memory usage for cached results for applying a route map.
Filter-list Cache Entries	Number of entries and memory usage for cached results for applying an AS path filter list.
Next Hop Cache Entries	Number of entries and memory usage for cached next-hop information.
Update messages queued	Number of update messages queued to be sent across all neighbors for which the specified process is responsible.
Read	Indicates the number of messages read by the process, the total size of read messages, the number of read operations performed, and the real time spent by the process performing read operations.
Read throttled	Number of times that reading from TCP has been throttled due to a backlog of messages read but not processed.
inbound messages	Number of read messages that have been processed and the real time spent processing inbound messages.
Wrote	Amount of data that has been written by the process, the number of write operations performed, and the real time spent by the process performing write operations.
Processing write list	Number of messages written from write lists, the number of times the write list has been processed, and the real time spent processing the write list. **Note** Write lists typically contain only update messages.
Processing write queue	Number of messages written from write queues, number of times the write queue has been processed, and the real time spent processing the write queue.
Socket setup	Number of socket setup operations performed and the real time spent during socket setup operations.
Configuration	Number of configuration requests received by the process and the real time spent processing configuration requests.
Operational data	Number of requests for operational data (for **show** commands) received by the process and the real time spent processing operation data requests

Field	Description
State	BGP system state for the specified address family and process. This may be one of the following: read-only mode—Initial set of updates is being recovered. In this mode, route selection is not performed, routes are not installed in the global RIB, and updates are not advertised to peers. best-path calculation mode—Route selection is being performed for the routes that were received while in read-only mode. import mode—Routes are imported from one VRF to another VRF once the best paths are calculated. This mode is supported in VPNv4 unicast address family mode. RIB update mode—Routes that were selected in best-path calculation mode are being installed in the global RIB. label allocation mode: Labels are allocated for the received prefixes based on the requirement. normal mode—Best paths are sent to the peers for routes that exist in the RIB. The route selection, import processing, RIB updates, and label allocation are performed as new updates are received.
BGP Table Version	Last version used in the BGP database for received routes.
Network Entries	Number of sets of prefix information held in the specified BGP process for the specified address family.
Soft Reconfig Entries	Number of sets of prefix information that are present only for the purpose of supporting soft reconfiguration.
Dampened Paths	Number of routes that are suppressed due to dampening for the specified address family.
History Paths	Number of routes that are currently withdrawn, but are being maintained to preserve dampening information.
Prefixes (Allocated/Freed)	Number of sets of prefix information for the specified address family that have been allocated and freed during the lifetime of the process.
Paths (Allocated/Freed)	Number of sets of route information for the specified address family that have been allocated and freed during the lifetime of the process.
Prefixes (Number/Memory Used)	Number of sets of prefix information currently allocated for the specified address family and amount of memory used by them.
Paths (Number/Memory Used)	Number of sets of route information currently allocated for the specified address family and amount of memory used by them.
Updates generated	Number of prefixes for which updates have been generated, the number of messages used to advertise the updates, the number of update generation runs performed, and the real time spent generating updates for the specified address family.

Field	Description
Scanner	Number of times the scanner has run for the specified address family and real time spent in scanner processing.
RIB Update	Number of global routing information base update runs performed for the specified address family, number of prefixes installed, withdrawn, or modified in the global RIB during these runs, and real time spent performing these runs.
Process has converged	Indicates whether the process has reached initial convergence for the specified address family.
First neighbor established	Time stamp (in seconds) recording the time the first neighbor in the process was established.
Entered DO_BESTPATH mode	Time stamp (in seconds) recording the time best-path calculation mode was entered.
Entered DO_RIBUPD mode	Time stamp (in seconds) recording the time RIB update mode was entered.
Entered Normal mode	Time stamp (in seconds) recording the time normal mode was entered.
Last UPDATE sent	Time stamp (in seconds) recording the time the last update was sent to a neighbor.

The following is sample output from the **show bgp vpnv4 unicast process performance-statistics detail** command:

```
RP/0/RP0/CPU0:router# show bgp vpnv4 unicast process performance-statistics detail
BGP Speaker process: 0, Node: node0_8_CPU0 Restart count: 1
                          Total           Nbrs Estab/Cfg
Default VRFs:             1               4/12
Non-Default VRFs:         1009            1082/1337

                          Sent            Received
Updates:                  362259          5688505
Notifications:            14              0

                          Number          Memory Used
Attributes:               14896           2979200
AS Paths:                 17              1100
Communities:              3               120
Extended communities:     1849            124440
Route Reflector Entries:  417             25020
Nexthop Entries:          2941            539572
Update messages queued:   0

                          Alloc           Free
Pool 210:                 28955629        28955628
Pool 310:                 363103          363103
Pool 600:                 4931162         4931162
Pool 1100:                104693          104693
Pool 4300:                799374          799374

Read 34755745 messages (3542094326 bytes) in 30528983 calls (time spent: 6427.769 secs)
Read partly throttled 1506 times
    Read 14 times after crossing lower threshold Processed 5836892 inbound update messages
 (time spent: 6229.512 secs)
```

```
 Wrote 825719955 bytes in 29272669 calls (time spent: 2318.472 secs)
 Processing sub-group: wrote 861402 messages in 1113810 calls (time spent: 145.446 secs)
 Processing write queue: wrote 6288 messages in 20498 calls (time spent: 0.039 secs)
 Socket setup (LPTS): 0 calls (time spent: 0.000 secs)
 event_file_attach calls: Input 8769, Output 2810, Input-output 0
Configuration: 989 requests (time spent: 0.046 secs) Operational data: 92396 requests (time
 spent: 98.864 secs)
Current Clock Time: not set Update Generation master timer:
    id: 0, time left: 0.0 sec, last processed: not set
    expiry time of parent node: not set
IO master timer:
    id: 0, time left: 0.0 sec, last processed: not set
    expiry time of parent node: not set


Address Family: VPNv4 Unicast
State: Normal mode.
BGP Table Version: 23211188
Attribute download: Disabled
Soft Reconfig Entries: 0

                   Last 8 Triggers       Ver         Tbl Ver

Label Thread       Jun 18 05:31:39.120   23211188    23211188
                   Jun 18 05:31:35.274   23211188    23211188
                   Jun 18 05:31:34.340   23211187    23211188
                   Jun 18 05:31:34.189   23211186    23211187
                   Jun 18 05:31:29.120   23211186    23211186
                   Jun 18 05:31:28.861   23211186    23211186
                   Jun 18 05:31:19.640   23211186    23211186
                   Jun 18 05:31:19.272   23211186    23211186
                   Total triggers: 639526

Import Thread      Jun 18 05:31:39.120   23211188    23211188
                   Jun 18 05:31:35.274   23211188    23211188
                   Jun 18 05:31:34.340   23211187    23211188
                   Jun 18 05:31:34.189   23211186    23211187
                   Jun 18 05:31:29.120   23211186    23211186
                   Jun 18 05:31:28.861   23211186    23211186
                   Jun 18 05:31:19.640   23211186    23211186
                   Jun 18 05:31:19.272   23211186    23211186
                   Total triggers: 689177

RIB Thread         Jun 18 05:31:39.146   23211188    23211188
                   Jun 18 05:31:35.299   23211188    23211188
                   Jun 18 05:31:34.525   23211187    23211188
                   Jun 18 05:31:34.494   23211186    23211188
                   Jun 18 05:31:34.340   23211186    23211188
                   Jun 18 05:31:34.255   23211186    23211188
                   Jun 18 05:31:29.146   23211186    23211186
                   Jun 18 05:31:28.886   23211186    23211186
                   Total triggers: 668084

Update Thread      Jun 18 05:31:39.171   ---         23211188
                   Jun 18 05:31:35.324   ---         23211188
                   Jun 18 05:31:34.558   ---         23211188
                   Jun 18 05:31:34.521   ---         23211188
                   Jun 18 05:31:34.327   ---         23211188
                   Jun 18 05:31:29.170   ---         23211186
                   Jun 18 05:31:28.910   ---         23211186
                   Jun 18 05:31:19.690   ---         23211186
                   Total triggers: 660143

                      Allocated       Freed
```

```
Remote Prefixes:      3150972         2885064
Remote Paths:         7639074         7118286

Local Prefixes:       3760870         3425614
Local Paths:          7892100         7595657

                      Number          Mem Used
Remote Prefixes:      265908          29781696
Remote Paths:         520788          24997824
Remote RDs:           12424           2832672

Local Prefixes:       335256          37548672
Local Paths:          296443          14229264
Local RDs:            1009            230052

Total Prefixes:       601164          67330368
Total Paths:          817231          39227088
Imported Paths:       265675          12752400
Total RDs:            13433           3062724
Same RDs:             0               0


Update Groups: 3  Subgroups: 2
Updates generated: 1438448 prefixes in 67375 messages from 181564 calls (time spent: 6779.576
 secs)
Scanner: 0 scanner runs (time spent: 0.000 secs) RIB update: 0 rib update runs, 0 prefixes
 installed, 0 modified,
            0 prefixes removed (time spent: 0.000 secs) RIB table update: 0 table deletes,
 0 table invalid, 3526736604 table skip,
            0 no local label, 0 rib retries Process has not converged for VPNv4 Unicast.

First neighbor established:  Jun 11 08:32:10
Entered DO_BESTPATH mode:    Jun 11 08:52:10
Entered DO_IMPORT mode:      Jun 11 08:52:12
Entered DO_LABEL_ALLOC mode: Jun 11 08:52:16
Entered DO_RIBUPD mode:      Jun 11 08:52:19
Entered Normal mode:         Jun 11 08:52:23
Latest UPDATE sent:          Jun 18 05:31:34
```

The following is sample output from show bgp process detail command with information on additional paths send and receive information:

```
BGP Process Information:
BGP is operating in STANDALONE mode
Autonomous System number format: ASDOT
Autonomous System: 100
Router ID: 22.22.22.22 (manually configured)
Default Cluster ID: 2.2.2.2 (manually configured)
Active Cluster IDs:  2.2.2.2
Fast external fallover enabled
Neighbor logging is enabled
Enforce first AS enabled
AS Path multipath-relax is enabled
Default local preference: 100
Default keepalive: 60
Graceful restart enabled
Restart time: 120
Stale path timeout time: 360
RIB purge timeout time: 600
Non-stop routing is enabled
Update delay: 120
Generic scan interval: 60
```

```
.....
.....
                           Allocated       Freed
Prefixes:                  12              0
Paths:                     60              0
Path-elems:                12              0

                           Number          Mem Used
Prefixes:                  12              1200
Paths:                     60              3120
Path-elems:                12              624
```

On Cisco 5500 series routers, the **Update wait-install enabled** section in the **show bgp process** command displays the delay of the BGP process update since the last router reload.

```
Router# show bgp process
Wed Aug 24 00:40:48.649 PDT

BGP Process Information:
BGP is operating in STANDALONE mode
Autonomous System number format: ASPLAIN
Autonomous System: 100
Router ID: 192.168.0.2 (manually configured)
Default Cluster ID: 192.168.0.2
Active Cluster IDs:  192.168.0.2
------------------------------
------------------------------
Update wait-install enabled:
  ack request 2, ack rcvd 2, slow ack 0
  startup delay 10 secs

--More—
```

show bgp regexp

To display routes matching the autonomous system path regular expression, use the**show bgp regexp** command in XR EXEC mode.

show bgp regexp *regular-expression*

Syntax Description

ipv4	(Optional) Specifies IP Version 4 address prefixes.
unicast	(Optional) Specifies unicast address prefixes.
multicast	(Optional) Specifies multicast address prefixes.
labeled-unicast	(Optional) Specifies labeled unicast address prefixes.
all	(Optional) For subaddress families, specifies prefixes for all subaddress families.
tunnel	(Optional) Specifies tunnel address prefixes.
mdt	(Optional) Specifies multicast distribution tree (MDT) address prefixes.
ipv6	(Optional) Specifies IP Version 6 address prefixes.
all	(Optional) For address family, specifies prefixes for all address families.
vpnv4 unicast	(Optional) Specifies VPNv4 unicast address families.
vrf	(Optional) Specifies VPN routing and forwarding (VRF) instance.
vrf-name	(Optional) Name of a VRF.
all	(Optional) For VRF, specifies all VRFs.
ipv4 { **unicast** | **labeled-unicast** }	(Optional) For VRF, specifies IPv4 unicast or labeled-unicast address families.
ipv6 unicast	(Optional) For VRF, specifies IPv6 unicast address families.
vpnv6 unicast	(Optional) Specifies VPNv6 unicast address families.
regular-expression	Regular expression to match the BGP autonomous system paths.

Command Default

If no address family or subaddress family is specified, the default address family and subaddress family specified using the **set default-afi** and **set default-safi** commands are used.

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

Note The **set default-afi** command is used to specify the default address family for the session, and the **set default-safi** command is used to specify the default subaddress family for the session. See the for detailed information and syntax for the **set default-afi** and **set default-safi** commands. If you do not specify a default address family, the default address family is IPv4. If you do not specify a default subaddress family, the default subaddress family is unicast.

BGP contains a separate routing table for each configured address family and subaddress family combination. The address family and subaddress family options specify the routing table to be examined. If the **all** keyword is specified for the address family or subaddress family, each matching routing table is examined in turn.

Use the **show bgp regexp** command to display all routes in the specified BGP table whose autonomous system path is matched by the specified regular expression.

Note If the regular expression contains spaces and parentheses, it must be specified and surrounded by quotation marks.

Task ID

Task ID	Operations
bgp	read

Examples

The following is sample output from the **show bgp regexp** command:

```
RP/0/RP0/CPU0:router# show bgp regexp "^3 "

BGP router identifier 10.0.0.5, local AS number 1
BGP main routing table version 64
BGP scan interval 60 secs
Status codes: s suppressed, d damped, h history, * valid, > best
              i - internal, S stale
Origin codes: i - IGP, e - EGP, ? - incomplete
   Network      Next Hop            Metric LocPrf Weight Path
*>i172.20.17.121      10.0.101.2                    100      0 3 2000 3000 i
*>i10.0.0.0           10.0.101.2                    100      0 3 100 1000 i
*>i172.5.23.0/24      10.0.101.2                    100      0 3 4 60 4378 i
```

This table describes the significant fields shown in the display.

Table 31: show bgp regexp Field Descriptions

Field	Description
BGP router identifier	BGP identifier for the local system.
local AS number	Autonomous system number for the local system.
Dampening enabled	Displayed if dampening has been enabled for the routes in this BGP routing table.

Field	Description
BGP main routing table version	Last version of the BGP database that was installed into the main routing table.
BGP scan interval	Interval (in seconds) between scans of the BGP table specified by the address family and subaddress family.
Status codes	Status of the table entry. The status is displayed as a three-character field at the beginning of each line in the table. The first character may be (in order of precedence): S—Path is stale, indicating that a graceful restart is in progress with the peer from which the route was learned. s—Path is more specific than a locally sourced aggregate route and has been suppressed. *—Path is valid. The second character may be (in order of precedence): >—Path is the best path to use for that network. d—Path is dampened. h—Path is a history entry, representing a route that is currently withdrawn, but that is being maintained to preserve dampening information. Such routes should never be marked as valid. The third character may be: i—Path was learned by an internal BGP (iBGP) session.
Origin codes	Origin of the path. The origin code is placed at the end of each line in the table. It can be one of the following values: i—Path originated from an Interior Gateway Protocol (IGP) and was advertised with a **network** or **aggregate-address** command. e—Path originated from an Exterior Gateway Protocol (EGP). ?—Origin of the path is not clear. Usually, this is a route that is redistributed into BGP from an IGP.
Network	IP address of a network entity.
Next Hop	IP address of the next system that is used when a packet is forwarded to the destination network. An entry of 0.0.0.0 indicates that the router has a non-BGP route to this network.
Metric	Value of the interautonomous system metric, otherwise known as the Multi Exit Discriminator (MED) metric.
LocPrf	Local preference value. This is used to determine the preferred exit point from the local autonomous system. It is propagated throughout the local autonomous system.
Weight	Path weight. Weight is used in choosing the preferred path to a route. It is not advertised to any neighbor.

Field	Description
Path	Autonomous system path to the destination network. At the end of the path is the origin code for the path.

show bgp route-policy

To display Border Gateway Protocol (BGP) information about networks that match an outbound route policy, use the **show bgp route-policy** command in XR EXEC mode.

show bgp route-policy *route-policy-name* [**standby**]

Syntax Description

ipv4	(Optional) Specifies IP Version 4 address prefixes.
unicast	(Optional) Specifies unicast address prefixes.
multicast	(Optional) Specifies multicast address prefixes.
labeled-unicast	(Optional) Specifies labeled unicast address prefixes.
all	(Optional) For subaddress families, specifies prefixes for all subaddress families.
tunnel	(Optional) Specifies tunnel address prefixes.
ipv6	(Optional) Specifies IP Version 6 address prefixes.
all	(Optional) For address family, specifies prefixes for all address families.
vpnv4 unicast	(Optional) Specifies VPNv4 unicast address families.
rd *rd-address*	(Optional) Displays routes with a specific route distinguisher.
vrf	(Optional) Specifies VPN routing and forwarding (VRF) instance.
vrf-name	(Optional) Name of a VRF.
all	(Optional) For VRF, specifies all VRFs.
ipv4 { **unicast** | **labeled-unicast** }	(Optional) For VRF, specifies IPv4 unicast or labeled-unicast address families.
ipv6 unicast	(Optional) For VRF, specifies IPv6 unicast address families.
route-policy-name	Name of a route policy.

Command Default

If no address family or subaddress family is specified, the default address family and subaddress family specified using the **set default-afi** and **set default-safi** commands are used.

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines



Note The **set default-afi** command is used to specify the default address family for the session, and the **set default-safi** command is used to specify the default subaddress family for the session. See the for detailed information and syntax for the **set default-afi** and **set default-safi** commands. If you do not specify a default address family, the default address family is IPv4. If you do not specify a default subaddress family, the default subaddress family is unicast.

BGP contains a separate routing table for each address family and subaddress family combination that has been configured. The address family and subaddress family options specify the routing table to be examined. If the **all** keyword is specified for the address family or subaddress family, each matching routing table is examined.

A route policy must be configured to use this command. When the **show bgp route-policy** command is entered, routes in the specified BGP table are compared with the specified route policy, and all routes passed by the route policy are displayed.

If a pass clause is encountered while the route policy is being applied to the route and the route policy processing completes without hitting a drop clause, the route is displayed. The route is not displayed if a drop clause is encountered, if the route policy processing completes without hitting a pass clause, or if the specified route policy does not exist.

The information displayed does not reflect modifications the policy might make to the route. To display such modifications, use the **show bgp policy** command.

Task ID

Task ID	Operations
bgp	read

Examples

The following is sample output from the **show bgp route-policy** command in XR EXEC mode:

```
RP/0/RP0/CPU0:router# show bgp route-policy p1

BGP router identifier 172.20.1.1, local AS number 1820
BGP main routing table version 729
Dampening enabled
BGP scan interval 60 secs
Status codes: s suppressed, d damped, h history, * valid, > best
              i - internal, S stale
Origin codes: i - IGP, e - EGP, ? - incomplete
   Network            Next Hop            Metric LocPrf Weight Path
*  10.13.0.0/16       192.168.40.24                          0 1878 704 701 200 ?
*  10.16.0.0/16       192.168.40.24                          0 1878 704 701 i
```

This table describes the significant fields shown in the display.

Table 32: show bgp route-policy Field Descriptions

Field	Description
BGP router identifier	BGP identifier for the local system.

Field	Description
local AS number	Autonomous system number for the local system.
BGP main routing table version	Last version of the BGP database that was installed into the main routing table.
Dampening enabled	Displayed if dampening is enabled for the routes in this BGP routing table.
BGP scan interval	Interval (in seconds) between scans of the BGP table specified by the address family and subaddress family.
Status codes	Status of the table entry. The status is displayed as a three-character field at the beginning of each line in the table. The first character may be (in order of precedence): S—Path is stale, indicating that a graceful restart is in progress with the peer from which the route was learned. s—Path is more specific than a locally sourced aggregate route and has been suppressed. *—Path is valid. The second character may be (in order of precedence): >—Path is the best path to use for that network. d—Path is dampened. h—Path is a history entry, representing a route that is currently withdrawn, but that is being maintained to preserve dampening information. Such routes should never be marked as valid. The third character may be: i—Path was learned by an internal BGP (iBGP) session.
Origin codes	Origin of the path. The origin code is displayed at the end of each line in the table. It can be one of the following values: i—Path originated from an Interior Gateway Protocol (IGP) and was advertised with a **network** or **aggregate-address** command. e—Path originated from an Exterior Gateway Protocol (EGP). ?—Origin of the path is not clear. Usually, this is a route that is redistributed into BGP from an IGP.
Network	IP prefix and prefix length for a network.
Next Hop	IP address of the next system that is used when a packet is forwarded to the destination network. An entry of 0.0.0.0 indicates that the router has a non-BGP route to this network.
Metric	Value of the interautonomous system metric, otherwise known as the Multi Exit discriminator (MED) metric.

Field	Description
LocPrf	Local preference value. This is used to determine the preferred exit point from the local autonomous system. It is propagated throughout the local autonomous system.
Weight	Path weight. Weight is used in choosing the preferred path to a route. It is not advertised to any neighbor.
Path	Autonomous system path to the destination network. At the end of the path is the origin code for the path.

show bgp ipv4 rt-filter neighbor

To verify the RT-constraint default-prefix state for a specified IPv4 neighbor, use the **show bgp ipv4 rt-filter neighbor** command in XR EXEC mode.

show bgp ipv4 rt-filter neighbor *ip-address*

Syntax Description

ip-address	Specifies the IPv4 address of the BGP neighbor.

Command Default

No default behavior or values

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 26.3.1	The command output was enhanced to display the RT-constraint default prefix state.

Usage Guidelines

No specific guidelines impact the use of this command.

Task ID

Task ID	Operations
bgp	read

Examples

This example verifies the RT-constraint default prefix state for neighbor `10.10.10.1`.

```
Router# show bgp ipv4 rt-filter neighbor 10.10.10.1 | include "RT-Constraint Default"
RT-Constraint Default Prefix: Suppressed-By-Config
```

No locally originated `0:0x0000:0:0/0` path is present.

This example verifies the RT-filter table.

```
Router# show bgp ipv4 rt-filter

Network          Next Hop Metric LocPrf Weight Path
*>i0:0x0000:0:0/0 11.11.11.3          100      0 i
Processed 1 prefixes, 1 paths
```

show bgp session-group

To display information about the Border Gateway Protocol (BGP) configuration for session groups, use the **show bgp session-group** command in XR EXEC mode.

show bgp session-group *group-name* {**configuration** [**defaults**] [**nvgen**] | **inheritance** | **users**}

Syntax Description

group-name	Name of the session family group to display.
configuration	(Optional) Displays the effective configuration for the session group, including any inherited configuration.
defaults	(Optional) Displays all configuration, including default configuration.
nvgen	(Optional) Displays output in the form of the **show running-config** command. If the **defaults** keyword also is specified, the output is not suitable for cutting and pasting into a configuration session.
inheritance	(Optional) Displays the session groups from which this session group inherits configuration.
users	(Optional) Display the session groups, neighbor groups, and neighbors that inherit configuration from this session group.

Command Default

No default behavior or value

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

Use the **show bgp session-group** command with the *group-name* **configuration** argument and keyword to display the effective configuration of a session group, including any configuration inherited from other session groups through application of the **use** command. The source for each configured command is also displayed.

Use the **defaults** keyword to display the value of all configuration, including default configuration. Use the **nvgen** keyword to display configuration in the form of the **show running-config** command output. Output in this form is suitable for cutting and pasting into a configuration session.

Use the **show bgp session-group** command with the *group-name* **inheritance** argument and keyword to display the session groups from which the specified session group inherits configuration.

Use the **show bgp session-group** command with the *group-name* **users** argument and keyword to display the neighbors, neighbor groups, and session groups that inherit configuration from the specified session group.

Task ID

Task ID	Operations
bgp	read

Examples

For the example shown here, the following configuration is used:

```
session-group group3
  advertisement-interval 5
  dmzlink-bw
 !
 session-group group1
  use session-group group2
  update-source Loopback0
 !
 session-group group2
  use session-group group3
  ebgp-multihop 2
```

The following example shows the **show bgp session-group** command with the **configuration** keyword:

```
RP/0/RP0/CPU0:router# show bgp session-group group1 configuration

session-group group1
 advertisement-interval 5[s:group2 s:group3]
 ebgp-multihop 2         [s:group2]
 update-source Loopback0 []
 dmzlink-bandwidth              [s:group2 s:group3]
```

The source of each command is shown to the right of the command. For example, **update-source** is configured directly on session group group1. The **dmzlink-bandwidth** command is inherited from session group group2, which in turn inherits it from session group group3.

The following example shows the **show bgp session-group** command with the **users** keyword:

```
RP/0/RP0/CPU0:router# show bgp session-group group2 users

IPv4 Unicast:a:group1
```

The following example shows the **show bgp session-group** command with the **inheritance** keyword.

```
RP/0/RP0/CPU0:router# show bgp session-group group1 inheritance

Session:s:group2 s:group3
```

The command output shows that the session group group1 directly uses the group2 session group. The group2 session group uses the group3 session group.

This table describes the significant fields shown in the display.

Table 33: show bgp session-group Field Descriptions

Field	Description
[ ]	Configures the command directly on the specified session group.
s:	Indicates the name that follows is a session group.
a:	Indicates the name that follows is an address family group.

Field	Description
n:	Indicates the name that follows is a neighbor group.
[dflt]	Indicates the command is not explicitly configured or inherited, and the default value for the command is used. This field may be shown when the **defaults** keyword is specified.
<not set>	Indicates that the default is for the command to be disabled. This field may be shown when the **defaults** keyword is specified.

show bgp sessions

To display brief information about BGP neighbors, use the **show bgp sessions** command in XR EXEC mode.

show bgp sessions [**not-established**] [**not-nsr-ready**] [**bfd-state**]

Syntax Description

not-established	(Optional) Displays all the neighbors that are not in established state
not-nsr-ready	(Optional) Displays all the neighbors that are not nonstop routing (NSR) ready.
bfd-state	(Optional) Displays BFD state information.

Command Default

No default behavior or values

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 6.0	This command was introduced.
Release 24.3.1	This command was modified. The **bfd-state** keyword was added. Also, the fields: BFDMode and BFDState were added to the output.

Usage Guidelines

The **show bgp sessions** command without a keyword provides brief information about all the BGP neighbors configured irrespective of the address family or VRF.

The **show bgp sessions** command with the **not-established** keyword shows BGP peers which are yet to establish their peering relationship.

The **show bgp session** command with the and **not-nsr-ready** keyword shows BGP peers which are yet to reach the nsr ready state.

Task ID

Task ID	Operations
bgp	read

Examples

The following is sample output from the **show bgp sessions** command in XR EXEC mode:

```
Router# show bgp sessions
Thu Jan 15 17:41:45.277 UTC

Neighbor        VRF                 Spk    AS   InQ  OutQ  NBRState     NSRState
10.2.2.2         default              0     1     0     0  Active        None
10.0.101.1      default               0     1     0     0  Established  NSR Ready
10.0.101.2      default               0     1     0     0  Established  NSR Ready
10.0.101.3      default               0     1     0     0  Established  NSR Ready
10.0.101.4      default               0     1     0     0  Established  NSR Ready
10.0.101.5      default               0     1     0     0  Established  NSR Ready
10.0.101.6      default               0     1     0     0  Established  NSR Ready
10.0.101.7      default               0     1     0     0  Established  NSR Ready
```

```
10.0.101.8      default                 0     1     0     0  Established  NSR Ready
10.0.101.9      default                 0     1     0     0  Established  NSR Ready
10.11.12.2      default                 0   100     0     0  Established  NSR Ready
10.0.0.2        900                     0     2     0     0  Established  NSR Ready
9000::1001      900                     0     2     0     0  Established  NSR Ready
91.0.0.2        901                     0     2     0     0  Established  NSR Ready
9100::1001      901                     0     2     0     0  Established  NSR Ready
92.0.0.2        902                     0     2     0     0  Established  NSR Ready
9200::1001      902                     0     2     0     0  Established  NSR Ready
93.0.0.2        903                     0     2     0     0  Established  NSR Ready
9300::1001      903                     0     2     0     0  Established  NSR Ready
94.0.0.2        904                     0     2     0     0  Established  NSR Ready
9400::1001      904                     0     2     0     0  Established  NSR Ready
95.0.0.2        905                     0     2     0     0  Established  NSR Ready
9500::1001      905                     0     2     0     0  Established  NSR Ready
96.0.0.2        906                     0     2     0     0  Established  NSR Ready
9600::1001      906                     0     2     0     0  Established  NSR Ready
97.0.0.2        907                     0     2     0     0  Established  NSR Ready
9700::1001      907                     0     2     0     0  Established  NSR Ready
98.0.0.2        908                     0     2     0     0  Established  NSR Ready
9800::1001      908                     0     2     0     0  Established  NSR Ready
99.0.0.2        909                     0     2     0     0  Idle         None
9900::1001      909                     0     2     0     0  Idle         None
12.13.14.16     red                     0     2     0     0  Idle         None
10.0.101.1      red                     0     2     0     0  Active       None
1234:5678:9876::1111
                red                     0     3     0     0  Idle         None
2020::1002      red                     0     2     0     0  Established  NSR Ready
10.2.3.4         this-is-a-long-vrf-name
                                        0     5     0     0  Idle         None
1111:2222:3333:4444:5555::6789
                this-is-a-long-vrf-name
                                        0     7     0     0  Idle         None
```

The following is sample output from the **show bgp sessions** command with the **not-established** keyword:

```
Router# show bgp sessions not-established
Fri Jan 30 11:30:42.720 PST PDT

Neighbor          VRF                   Spk    AS   InQ  OutQ  NBRState     NSRState
10.0.101.5        default     0       100     0       0       Active   None
192.0.2.1         vrf1_1      0       302     0       0       Idle     None
203.0.113.1       vrf1_1      0       302     0       0       Idle     None
203.1.113.1       vrf1_1      0       302     0       0       Idle     None
203.2.113.1       vrf1_1      0       302     0       0       Idle     None
209.165.200.225   vrf1_1      0       304     0       0       Idle     None
2008:2:2:2::2     vrf1_1      0       304     0       0       Idle     None
209.165.202.129   vrf1_1      0       302     0       0       Idle     None
```

The following is sample output from the **show bgp sessions** command with the **not-nsr-ready** keyword:

```
Router# show bgp sessions not-nsr-ready
Fri Jan 30 11:30:52.301 PST PDT

Neighbor          VRF                   Spk    AS   InQ  OutQ  NBRState     NSRState
10.0.101.5        default               0     100   0    0     Active       None
192.0.2.1         vrf1_1                0     302   0    0     Idle         None
203.0.113.1       vrf1_1                0     302   0    0     Idle         None
203.1.113.1       vrf1_1                0     302   0    0     Idle         None
203.2.113.1       vrf1_1                0     302   0    0     Idle         None
```

```
209.165.200.225     vrf1_1              0     304   0    0     Idle         None
2008:2:2:2::2       vrf1_1              0     302   0    0     Idle         None
209.165.202.129     vrf2_1              0     302   0    0     Idle         None
```

The following is sample output from the **show bgp sessions** command with the **bfd-state** keyword. The BFDMode column shows the operational mode of BFD for the BGP session. Different modes represent different behaviors for the BFD protocol. The BFDState column shows the current state of the BFD session.

```
Router# show bgp sessions bfd-state
Fri Jan 30 11:30:52.301 PST PDT

Neighbor      VRF      Spk  AS  InQ  OutQ  NBRState     NSRState  BFDMode           BFDState
10.1.1.1      default  0    1   0    0     Established  None      Disabled          BFD
state unknown
10.1.1.78     default  0    1   0    0     Established  None      StrictModeOverride Up
10.1.1.82     default  0    1   0    0     Idle         None      StrictModeOverride Down
10.1.1.86     default  0    1   0    0     Established  None      StrictMode         Up
10.1.1.94     default  0    1   0    0     Established  None      StrictMode         Down
10.1.1.90     default  0    1   0    0     Established  None      XRStrictMode       Up
10.1.1.98     default  0    1   0    0     Idle         None      XRStrictMode       Down
10.1.1.114    default  0    1   0    0     Established  None      Default            Down
```

This table describes the significant fields shown in the display.

Table 34: show bgp sessions Field Descriptions

Field	Description
Neighbor	Displays neighbor IP address.
VRF	Displays information about the VRF.
Spk	Speaker process that is responsible for the neighbor. Always 0.
AS	Autonomous system.
InQ	Number of messages from a neighbor waiting to be processed.
OutQ	Number of messages waiting to be sent to a neighbor.
NBRState	State of the Border Gateway Protocol (BGP) neighbor sessions.
NSRState	State of the Border Gateway Protocol (BGP) nonstop routing (NSR).
BFDMode	Operational mode of BFD for the BGP session. Different modes may represent different configurations or behaviors for the BFD protocol.
Disabled	BFD is not enabled on the BGP session. No BFD packets are sent or received, and BFD is not used to detect link failures for this session.
StrictMode	The device holds down the neighbor session until BFD session is up. This behaviour is in compliance with the IETF standards.
StrictModeOverride	BFD enforces the IETF strict mode for BFD, regardless of the capabilities negotiated between the two peers.

Field	Description
XRStrictMode	BFD prevents the establishment of BGP neighbor session until the BFD session is also up and running. This behaviour is in compliance with Cisco's propriety definition of BFD.
Default	BFD operates in its default mode, which means it uses the standard settings for detection time and other parameters. This mode uses the default values specified by the BFD protocol.
BFDState	Current state of the BFD session.
BFD state unknown	BFD session status is not known.
Up	BFD session is up and operational and BFD protocol has not detected any faults in the path to the neighbor.
Down	BFD session is down.

show bgp summary

To display the status of all Border Gateway Protocol (BGP) connections, use the **show bgp summary** command in XR EXEC mode.

show bgp [**ipv4** {**unicast** | **multicast** | **labeled-unicast** | **all** | **tunnel** | **mdt**} | **ipv6** {**unicast** | **multicast** | **all** | **labeled-unicast**} | **all** {**unicast** | **multicast** | **all** | **labeled-unicast** | **mdt** | **tunnel**} | **vpnv4 unicast** | **vrf** {**vrf-name** | **all**} [**ipv4** {**unicast** | **labeled-unicast**} | **ipv6 unicast**] | **vpvn6 unicast**] **summary**

Syntax Description

ipv4	(Optional) Specifies IP Version 4 address prefixes.
unicast	(Optional) Specifies unicast address prefixes.
multicast	(Optional) Specifies multicast address prefixes.
labeled-unicast	(Optional) Specifies labeled unicast address prefixes.
all	(Optional) For subaddress families, specifies prefixes for all subaddress families.
tunnel	(Optional) Specifies tunnel address prefixes.
multicast	(Optional) Specifies multicast address prefixes.
ipv6	(Optional) Specifies IP Version 6 address prefixes.
all	(Optional) For address family, specifies prefixes for all address families.
vpnv4 unicast	(Optional) Specifies VPNv4 unicast address families.
vrf	(Optional) Specifies VPN routing and forwarding (VRF) instance.
vrf-name	(Optional) Name of a VRF.
all	(Optional) For VRF, specifies all VRFs.
ipv4 { **unicast** | **labeled-unicast** }	(Optional) For VRF, specifies IPv4 unicast or labeled-unicast address families.
ipv6 unicast	(Optional) For VRF, specifies IPv6 unicast address families.

Command Default

If no address family or subaddress family is specified, the default address family and subaddress family specified using the **set default-afi** and **set default-safi** commands are used.

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

The **set default-afi** command is used to specify the default address family for the session, and the **set default-safi** command is used to specify the default subaddress family for the session. See for detailed information and syntax for the **set default-afi** and **set default-safi** commands. If you do not specify a default address family, the default address family is IPv4. If you do not specify a default subaddress family, the default subaddress family is unicast.

Use the **show bgp summary** command to display a summary of the neighbors for which the specified address family and subaddress family are enabled. If the neighbor does not have the specified address family and subaddress family enabled, it is not included in the output of the **show** command. If the **all** keyword is specified for the address family or subaddress family, a summary for each combination of address family and subaddress family is displayed in turn.

The table versions shown in the output (RcvTblVer, bRIB/RIB, SendTblVer, and TblVer) are specific to the specified address family and subaddress family. All other information is global.

The table versions provide an indication of whether BGP is up to date with all work for the specified address family and subaddress family.

- bRIB/RIB < RecvTblVer—Some received routes have not yet been considered for installation in the global routing table.
- TblVer < SendTblVer—Some received routes have been installed in the global routing table but have not yet been considered for advertisement to this neighbor.

Task ID

Task ID	Operations
bgp	read

Examples

The following is sample output from the **show bgp summary** command:

```
RP/0/RP0/CPU0:router#show bgp summary

BGP router identifier 10.0.0.0, local AS number 2
BGP generic scan interval 60 secs
BGP table state: Active
Table ID: 0xe0000000
BGP main routing table version 1
BGP scan interval 60 secs

BGP is operating in STANDALONE mode.

Process         RecvTblVer    bRIB/RIB  LabelVer  ImportVer  SendTblVer
Speaker                  1           0         1          1           0

Neighbor        Spk    AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  St/PfxRcd
10.0.101.0        0     2       0       0        0    0    0 00:00:00 Idle
10.0.101.1        0     2       0       0        0    0    0 00:00:00 Idle
```

This table describes the significant fields shown in the display.

Table 35: show bgp summary Field Descriptions

Field	Description
BGP router identifier	IP address of the router.

Field	Description
local AS number	Autonomous system number set by the router bgp, on page 256 command. • Range for 2-byte Autonomous system numbers (ASNs) is 1 to 65535. • Range for 4-byte Autonomous system numbers (ASNs) in asplain format is 1 to 4294967295. • Range for 4-byte Autonomous system numbers (ASNs) is asdot format is 1.0 to 65535.65535.
BGP generic scan interval	Interval (in seconds) between scans of the BGP table by a generic scanner.
BGP table state	State of the BGP database.
Table ID	BGP database identifier.
BGP main routing table version	Last version of the BGP database that was injected into the main routing table.
Dampening enabled	Displayed if dampening has been enabled for the routes in this BGP routing table.
BGP scan interval	Interval (in seconds) between scans of the BGP table specified by the address family and subaddress family.
BGP is operating in	Specifies BGP is operating in standalone mode.
Process	BGP process.
RecvTblVer	Last version used in the BGP database for received routes.
bRIB/RIB	Last version of the local BGP database that was injected into the main routing table.
LabelVer	Label version used in the BGP database for label allocation.
ImportVer	Last version of the local BGP database for importing routes.
SendTblVer	Latest version of the local BGP database that is ready to be advertised to neighbors.
Some configured eBGP neighbors do not have any policy	Some external neighbors exist that do not have both an inbound and outbound policy configured for every address family, using the **route-policy (BGP)** command. In this case, no prefixes are accepted and advertised to those neighbors.
Neighbor	IP address of a neighbor.
Spr	Speaker process that is responsible for the neighbor. Always 0.
AS	Autonomous system.
MsgRcvd	Number of BGP messages received from a neighbor.
MsgSent	Number of BGP messages sent to a neighbor.

Field	Description
TblVer	Last version of the BGP database that was sent to a neighbor.
InQ	Number of messages from a neighbor waiting to be processed.
OutQ	Number of messages waiting to be sent to a neighbor.
Up/Down	Length of time in (hh:mm:ss) that the BGP session has been in Established state, or the time since the session left Established state, if it is not established.
St/PfxRcd	If the BGP session is not established, the current state of the session. If the session is established, the number of prefixes the router has received from the neighbor. If the number of prefixes received exceeds the maximum allowed (as set by the **maximum-prefix** command), "(PfxRcd)" appears. If the connection has been shut down using the **shutdown** command, "(Admin)" appears. If the neighbor is external and it does not have an inbound and outbound policy configured for every address family, an exclamation mark (!) is inserted at the end of the state when using the **route-policy** (**BGP**) command. If the connection has been shut down due to out of memory (OOM), "(OOM)" appears.

show bgp truncated-communities

To display routes in the Border Gateway Protocol (BGP) routing table for which inbound policy or aggregation has exceeded the maximum number of communities that may be attached, use the **show bgp truncated-communities** command in XR EXEC mode.

show bgptruncated-communities standby

Syntax Description

ipv4	(Optional) Specifies IP Version 4 address prefixes.
unicast	(Optional) Specifies unicast address prefixes.
multicast	(Optional) Specifies multicast address prefixes.
labeled-unicast	(Optional) Specifies labeled unicast address prefixes.
all	(Optional) For subaddress families, specifies prefixes for all subaddress families.
tunnel	(Optional) Specifies tunnel address prefixes.
mdt	(Optional) Specifies multicast distribution tree (MDT) address prefixes.
multicast	(Optional) Specifies multicast address prefixes.
ipv6	(Optional) Specifies IP Version 6 address prefixes.
all	(Optional) For address family, specifies prefixes for all address families.
vpnv4 unicast	(Optional) Specifies VPNv4 unicast address families.
rd *rd-address*	(Optional) Displays routes with a specific route distinguisher.
vrf	(Optional) Specifies VPN routing and forwarding (VRF) instance.
vrf-name	(Optional) Name of a VRF.
all	(Optional) For VRF, specifies all VRFs.
ipv4 { **unicast** | **labeled-unicast** }	(Optional) For VRF, specifies IPv4 unicast or labeled-unicast address families.
ipv6 unicast	(Optional) For VRF, specifies IPv6 unicast address families.
vpvn6 unicast	(Optional) Specifies VPNv6 unicast address families.
standby	(Optional) Displays information about the standby card.

Command Default

If no address family or subaddress family is specified, the default address family and subaddress family specified using the **set default-afi** and **set default-safi** commands are used.

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

The **set default-afi** command is used to specify the default address family for the session, and the **set default-safi** command is used to specify the default subaddress family for the session. See the for detailed information and syntax for the **set default-afi** and **set default-safi** commands. If you do not specify a default address family, the default address family is IPv4. If you do not specify a default subaddress family, the default subaddress family is unicast.

BGP contains a separate routing table for each address family and subaddress family combination that has been configured. The address family and subaddress family options specify the routing table to be examined. If the **all** keyword is specified for the address family or subaddress family, each matching routing table is examined.

Use the **show bgp truncated-communities** command to display those routes in the specified BGP routing table in which the buffers used to store communities or extended communities have overflowed. An overflow occurs if an attempt is made to associate more communities or extended communities with the route than fits in a BGP update message. This can happen due to modification of communities or extended communities during aggregration or when inbound policy is applied.

Task ID

Task ID	Operations
bgp	read

Examples

The following is sample output from the **show bgp truncated-communities** command:

```
RP/0/RP0/CPU0:router# show bgp truncated-communities

BGP router identifier 172.20.1.1, local AS number 1820
BGP main routing table version 3042
BGP scan interval 60 secs
Status codes: s suppressed, d damped, h history, * valid, > best
              i - internal, S stale
Origin codes: i - IGP, e - EGP, ? - incomplete
   Network            Next Hop            Metric LocPrf Weight Path
*  10.13.0.0/16       192.168.40.24                          0 1878 704 701 200 ?
*> 10.16.0.0/16       192.168.40.24                          0 1878 704 701 i
```

This table describes the significant fields shown in the display.

Table 36: show bgp truncated-communities Field Descriptions

Field	Description
BGP router identifier	BGP Identifier for the local system.
local AS number	Autonomous system number for the local system.
BGP main routing table version	Last version of the BGP database that was installed into the main routing table.

Field	Description
Dampening enabled	Displayed if dampening is enabled for the routes in this BGP routing table.
BGP scan interval	Interval (in seconds) between scans of the BGP table specified by the address family and subaddress family.
Status codes	Status of the table entry. The status is displayed as a three-character field at the beginning of each line in the table. The first character may be (in order of precedence): S—Path is stale, indicating that a graceful restart is in progress with the peer from which the route was learned. s—Path is more specific than a locally sourced aggregate route and has been suppressed. *—Path is valid. The second character may be (in order of precedence): >—Path is the best path to use for that network. d—Path is dampened. h—Path is a history entry, representing a route that is currently withdrawn, but that is being maintained to preserve dampening information. Such routes should never be marked as valid. The third character may be: i—Path was learned by an internal BGP (iBGP) session.
Origin codes	Origin of the path. The origin code is displayed at the end of each line in the table. It can be one of the following values: i—Path originated from an Interior Gateway Protocol (IGP) and was advertised with a **network** or **aggregate-address** command. e—Path originated from an Exterior Gateway Protocol (EGP). ?—Origin of the path is not clear. Usually, this is a route that is redistributed into BGP from an IGP.
Network	IP prefix and prefix length for a network.
Next Hop	IP address of the next system that is used when a packet is forwarded to the destination network. An entry of 0.0.0.0 indicates that the router has a non-BGP route to this network.
Metric	Value of the interautonomous system metric, otherwise known as the Multi Exit Discriminator (MED) metric.
LocPrf	Local preference value. This is used to determine the preferred exit point from the local autonomous system. It is propagated throughout the local autonomous system.
Weight	Path weight. Weight is used in choosing the preferred path to a route. It is not advertised to any neighbor.

Field	Description
Path	Autonomous system path to the destination network. At the end of the path is the origin code for the path.

show bgp update-group

To display Border Gateway Protocol (BGP) information for update groups, use the **show bgp update-group** command in XR EXEC mode.

show bgp [**ipv4** {**unicast** | **multicast** | **labeled-unicast** | **all** | **tunnel** | **mdt**} | **ipv6** {**unicast** | **multicast** | **all** | **labeled-unicast**} | **all** {**unicast** | **multicast** | **all** | **labeled-unicast** | **mdt** | **tunnel**} | **vpnv4 unicast** | **vrf** {**vrf-name** | **all**} [**ipv4** {**unicast** | **labeled-unicast**} | **ipv6 unicast**] | **vpnv6 unicast**] **update-group** [**neighbor ip-address** | **process-id . index** [**summary** | **performance-statistics**]]

Syntax Description

ipv4	(Optional) Specifies IP Version 4 update groups.
unicast	(Optional) Specifies unicast update groups.
multicast	(Optional) Specifies multicast update groups.
labeled-unicast	(Optional) Specifies labeled unicast address prefixes.
all	(Optional) Displays both unicast and multicast update groups.
tunnel	(Optional) Specifies tunnel address prefixes.
ipv6	(Optional) Specifies IP Version 6 update groups.
all	(Optional) Displays both IP Version 4 and IP Version 6 update groups.
vpnv4 unicast	(Optional) Specifies VPNv4 unicast address families.
rd *rd-address*	(Optional) Displays routes with a specific route distinguisher.
vrf	(Optional) Specifies VPN routing and forwarding (VRF) instance.
vrf-name	(Optional) Name of a VRF.
all	(Optional) For VRF, specifies all VRFs.
ipv4 { **unicast** | **labeled-unicast** }	(Optional) For VRF, specifies IPv4 unicast or labeled-unicast address families.
ipv6 unicast	(Optional) For VRF, specifies IPv6 unicast address families.
neighbor *ip-address*	(Optional) Specifies information on an update group for a specific neighbor.
process-id.index	(Optional) Update group index. Process ID range is 0 to 254. Index range is 0 to 4294967295. The *process id.index* argument is specified as follows: process ID (dot) index. In standalone mode, the process ID is always 0.
summary	(Optional) Specifies summary of update group members.

performance-statistics	(Optional) Specifies performance information about the updates generated for the update group.

Command Default

If no address family or subaddress family is specified, the default address family and subaddress family specified using the **set default-afi** and **set default-safi** commands are used.

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

The **set default-afi** command is used to specify the default address family for the session, and the **set default-safi** command is used to specify the default subaddress family for the session. See the for detailed information and syntax for the **set default-afi** and **set default-safi** commands. If you do not specify a default address family, the default address family is IPv4. If you do not specify a default subaddress family, the default subaddress family is unicast.

Every BGP neighbor is automatically assigned to an update group for each address family that is enabled on the neighbor. Neighbors that have similar outbound policy, such that they are sent the same updates, are placed in the same update group.

Use the **show bgp update-group** command to display the update groups and a list of the neighbors that belong to the update group.

Use the **show bgp update-group neighbor** command to display details about the update group to which a neighbor belongs for the specified address family.

Use the **summary** keyword to display a summary of the neighbors belonging to the specified update group. The display format is the same as for the show bgp summary, on page 433 command.

Use the **performance-statistics** keyword to display information about the number of prefixes processed and the time taken to generate updates for the specified update group.

Note Update group indexes are not necessarily persistent over a process restart. If a BGP process restarts, the index of the update group to which a particular neighbor is assigned may be different, though the set of neighbors belonging to the update group is the same.

Task ID

Task ID	Operations
bgp	read

Examples

The following is sample output from the **show bgp update-group** command:

```
RP/0/RP0/CPU0:router# show bgp update-group

Update group for IPv4 Unicast, index 0.1:
  Attributes:
```

```
  Internal
  Common admin
  Send communities
  Send extended communities
  Minimum advertisement interval: 300
Update group desynchronized: 0
Sub-groups merged: 0
Messages formatted: 0, replicated: 0
Neighbors not in any sub-group:
  10.0.101.1
```

This table describes the significant fields shown in the display.

Table 37: show bgp update-group Field Descriptions

Field	Description
Update group for	Address family to which updates in this update group apply.
index	Update group index.
Attributes	Attributes common to all members of the update group.
Unsuppress map	Unsuppress route map used to selectively unsuppress more specific routes of locally generated aggregates for members of this update group.
Outbound policy	Route policy applied to outbound updates generated for members of this update group.
Internal	Members of the update group are internal peers.
ORF Receive enabled	Members of this update group are capable of receiving an outbound route filter.
Route Reflector Client	Local system is acting as a route reflector for members of this update group.
Remove private AS numbers	Members of this update group have private AS numbers stripped from outbound updates.
Next-hop-self enabled	Next- Next hop for members of the update group is set to the local router.
Directly connected IPv6 EBGP	Members of this update group are directly connected external BGP IPv6-based peers.
Configured Local AS	Local autonomous system (AS) used for members of this update group.
Common admin	Peers in this update group are under common administration (internal or confederation peers).
Send communities	Communities are sent to neighbors in this update group.
Send extended communities	Extended communities is sent to neighbors in this update group.
Minimum advertisement interval	Minimum advertisement interval for members of this update group.
replicated	Number of update messages replicated for this update group.

Field	Description
Messages formatted	Number of update messages generated for this update group.
Neighbors in this update group	List of neighbors that use this update group for the given address family.
Update group desynchronized	Number of times an update group has been split to accommodate the slower peer. This option is disabled.
Sub-groups merged	Number of times an update group has been split and merged.
Neighbors not in any sub-group	BGP neighbor that does not belong to any subgroup.

The following is sample output from the **show bgp update-group** command with the **ipv4** , **unicast** , and **summary** keywords and the *process id.index* argument:

```
RP/0/RP0/CPU0:router# show bgp ipv4 unicast update-group 0.1 summary

BGP router identifier 10.140.140.1, local AS number 1.1
BGP generic scan interval 60 secs
BGP table state: Active
Table ID: 0xe0000000
BGP main routing table version 1
BGP scan interval 60 secs

BGP is operating in STANDALONE mode.

Process         RecvTblVer    bRIB/RIB  LabelVer  ImportVer  SendTblVer
Speaker                  1           0         1          1           0

Neighbor        Spr    AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  St/PfxRcd

172.25.11.8       0     1       0       0        0    0    0 00:00:00 Idle
```

This is sample output from the **show bgp ipv4 unicast update-group** command showing the status of advertised permanent paths:

```
RP/0/RP0/CPU0:router# show bgp ipv4 unicast update-group
Update group for IPv4 Unicast, index 0.2:
  Attributes:
    Neighbor sessions are IPv4
    Outbound policy: PASS
    Internal
    Common admin
    First neighbor AS: 30813
    Send communities
    Send extended communities
    Next-hop-self enabled
    4-byte AS capable
    Non-labeled address-family capable
    Advertise Permanent-Network capable
    Send AIGP
    Minimum advertisement interval: 0 secs
  Update group desynchronized: 0
  Sub-groups merged: 4
  Number of refresh subgroups: 0
  Messages formatted: 42, replicated: 68
  Neighbors not in any sub-group:
```

```
100.12.13.3     100.13.13.3
```

This table describes the significant fields shown in the display.

Table 38: show bgp ipv4 unicast update-group Field Descriptions

Field	Description
BGP router identifier	IP address of the router.
local AS number	Autonomous system number set by the router bgp, on page 256 command. • Range for 2-byte Autonomous system numbers (ASNs) is 1 to 65535. • Range for 4-byte Autonomous system numbers (ASNs) in asplain format is 1 to 4294967295. • Range for 4-byte Autonomous system numbers (ASNs) is asdot format is 1.0 to 65535.65535.
BGP generic scan interval	Interval (in seconds) between scans of the BGP table by a generic scanner.
BGP table state	State of the BGP database.
Table ID	BGP database identifier.
BGP main routing table version	Last version of the BGP database that was injected into the main routing table.
Dampening enabled	Displayed if dampening has been enabled for the routes in this BGP routing table.
BGP scan interval	Interval (in seconds) between scans of the BGP table specified by the address family and subaddress family.
BGP is operating in	BGP is operating in standalone mode.
Process	BGP process.
RecvTblVer	Last version used in the BGP database for received routes.
bRIB/RIB	Last version of the local BGP database that was injected into the main routing table.
LabelVer	Label version used in the BGP database for label allocation.
ImportVer	Last version of the local BGP database for importing routes.
SendTblVer	Latest version of the local BGP database that is ready to be advertised to neighbors.
Some configured eBGP neighbors do not have any policy	Some external neighbors that exist do not have both an inbound and outbound policy configured for every address family, using the **route-policy (BGP)** command. In this case, no prefixes are accepted or advertised to those neighbors.
Neighbor	IP address of a neighbor.

Field	Description
Spr	Speaker process that is responsible for the neighbor. Always 0.
AS	Autonomous system.
MsgRcvd	Number of BGP messages received from a neighbor.
MsgSent	Number of BGP messages sent to a neighbor.
TblVer	Last version of the BGP database that was sent to a neighbor.
InQ	Number of messages from a neighbor waiting to be processed.
OutQ	Number of messages waiting to be sent to a neighbor.
Up/Down	Length of time (in hh:mm:s) that the BGP session has been in Established state, or the time since the session left Established state, if it is not established.
St/PfxRcd	If the BGP session is not established, the current state of the session. If the session is established, the number of prefixes the router has received from the neighbor. If the number of prefixes received exceeds the maximum allowed (as set by the **maximum-prefix** command), "(PfxRcd)" appears. If the connection has been shut down using the **shutdown** command, "(Admin)" appears. If the neighbor is external and it does not have an inbound and outbound policy configured for every address family, an exclamation mark (!) is inserted at the end of the state when using the **route-policy** (**BGP**) command.

show bgp vrf

To display Border Gateway Protocol (BGP) prefix information for VPN routing and forwarding (VRF) instances, use the **show bgp vrf** command in XR EXEC mode.

show bgp vrf { **all** *vrf-name* } **{ process** | **rpki** *refresh-list* } { **ipv4** { **unicast** [*ipv4-address/length* [**detail**]] | **labeled-unicast** } | **ipv6** { **unicast** } | **imported-routes** { **neighbor** | **standby** | **vrf** *vrf-name* }}

Syntax Description

vrf-name	Displays imported routes for a specific VRF.
all	Displays imported routes for all VRFs.
ipv4 { **unicast** | **labeled-unicast** }	(Optional) Specifies IP Version 4 unicast or labeled-unicast imported routes.
ipv6 unicast	(Optional) Specifies IP Version 6 unicast imported routes.
vrf *source-vrf-name*	(Optional) Displays routes imported from the specified source VRF.
neighbor *neighbor-address*	(Optional) Displays preview advertisements for a specified neighbor.
standby	(Optional) Displays information about the standby card.
process	Displays process information.
rpki *refresh-list*	Displays BGP Resource Public Key Infrastructure (RPKI) information.

Command Default

No default behavior or values

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 6.0	This command was introduced.
Release 6.1	This command was modified. The **process** keyword and *rpki refresh-list* keyword-argument combination was added.

Usage Guidelines

Use the **show bgp vrf imported-routes** command to display all paths imported into a specified VRF from the default VRF. Use the **neighbor** *neighbor-address* keyword and argument to display all imported paths and which paths were learned from the specified neighbor. Use the **vrf** *source-vrf-name* keyword and argument to display all imported routes that belong to the specified source VRF. The **neighbor** *neighbor-address* and **vrf** *source-vrf-name* cannot coexist.

Task ID

Task ID	Operations
bgp	read

Examples

The following is sample output from the **show bgp vrf imported-routes** command:

```
RP/0/RP0/CPU0:router# show bgp vrf vrf-1 ipv6 unicast imported-routes

BGP VRF one, state: Active BGP
BGP Route Distinguisher: 100:222
VRF ID: 0x60000001
BGP router identifier 10.2.0.1, local AS number 100
BGP table state: Active
Table ID: 0xe0800001
BGP main routing table version 41534

Status codes: s suppressed, d damped, h history, * valid, > best
              i - internal, S stale
Origin codes: i - IGP, e - EGP, ? - incomplete
   Network            Neighbor        Route Distinguisher    Source VRF
*>i1234:1052::/32     10.1.0.1        100:111                default
*>i2008:1:1:1::/112   10.1.0.1        100:111                default
*>i2008:111:1:1::1/128
                      10.1.0.1        100:111                default

Processed 3 prefixes, 3 paths
```

The following is sample output from the **show bgp vrf** *vrf-name* **ipv4 unicast** *ipv4-address/length* **detail** command.

```
RP/0/RP0/CPU0:router#  show bgp vrf foo ipv4 unicast 100.1.1.1/32 detail

Mon Dec  8 23:24:50.243 PST
BGP routing table entry for 100.1.1.1/32, Route Distinguisher:
30.30.30.30:0
Versions:
  Process           bRIB/RIB  SendTblVer
  Speaker                 43          43
    Local Label: 24001 (with rewrite);
    Flags: 0x05081001+0x00000200;
Last Modified: Dec  8 18:04:21.000 for 05:20:30
Paths: (1 available, best #1)
  Advertised to PE peers (in unique update groups):
    32.0.0.2
  Path #1: Received by speaker 0
  Flags: 0x400061000d060005, import: 0x80
  Advertised to PE peers (in unique update groups):
    32.0.0.2
  Local
    11.0.0.1 (metric 2) from 20.0.0.1 (11.0.0.1)
      Received Label 1234
      Origin IGP, localpref 100, valid, internal, best, group-best, import-candidate,
imported, reoriginated
      Received Path ID 0, Local Path ID 1, version 43
      Extended community: Encapsulation Type:8 Router MAC:aabb.ccdd.eeff RT:1:2
      Originator: 11.0.0.1, Cluster list: 20.20.20.20
      RIB RNH: table_id 0xe0000011, Encap 8, VNI 1234, MAC Address: aabb.ccdd.eeff, IP
Address: 11.0.0.1, IP table_id 0xe0000000
```

```
      Source AFI: L2VPN EVPN, Source VRF: default, Source Route
Distinguisher: 100:1
```

This table describes the significant fields shown in the display output for **show bgp vrf** command.

Table 39: show bgp vrf Field Descriptions

Field	Description
BGP VRF	VRF name.
state	State of the VRF.
BGP Route Distinguisher:	Unique identifier for the BGP routing instance.
VRF Id	VRF identifier.
BGP router identifier	IP address of the router.
local AS number	Autonomous system number set by the router bgp, on page 256 command. • Range for 2-byte Autonomous system numbers (ASNs) is 1 to 65535. • Range for 4-byte Autonomous system numbers (ASNs) in asplain format is 1 to 4294967295. • Range for 4-byte Autonomous system numbers (ASNs) is asdot format is 1.0 to 65535.65535.
BGP table state	State of the BGP database.
Table ID	Table identifier.
BGP main routing table version	Last version of the BGP database that was injected into the main routing table.
Network	Network address.
Neighbor	IP address of a neighbor.
Route Distinguisher	Unique identifier for the routing instance.
Source VRF	Source VRF for the imported route.

show protocols (BGP)

To display information about the Border Gateway Protocol (BGP) instances running on the router, use the **show protocols** command in XR EXEC mode and specify either the **bgp** or **all** keyword.

show protocols [**ipv4** | **ipv6** | **afi-all**] [**all***protocol*]

Syntax Description

ipv4	(Optional) Specifies the IP Version 4 address family.
ipv6	(Optional) Specifies the IP Version 6 address family.
afi-all	(Optional) Specifies all address families.
all	(Optional) Specifies all protocols for a given address family.
protocol	(Optional) Specifies a routing protocol. For the IPv4 address family, the options are **bgp** , **isis** , **rip** , and **ospf** . For the IPv6 address family, the options are **bgp** , **isis** , and **ospfv3** .

Command Default

Default is IPv4.

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

Use the **show protocols** command to get information about the protocols running on the router and to quickly determine which protocols are active. The command is designed to summarize the important characteristics of the running protocol, and command output varies depending on the specific protocol selected. For BGP, the command output lists the protocol ID, peers with elapsed time since last reset, and miscellaneous information, such as external and internal local distances and sourced routes.

Task ID

Task ID	Operations
bgp	read
rib	read

Examples

The following example shows the display for the **show protocols** command using the **bgp** keyword:

```
RP/0/RP0/CPU0:router# show protocols bgp

Routing Protocol "BGP 40"

Address Family IPv4 Unicast:
```

```
Distance: external 20 internal 200 local 200
Sourced Networks:
  10.100.0.0/16 backdoor
  10.100.1.0/24
  10.100.2.0/24
Routing Information Sources:
  Neighbor          State/Last update received
  10.5.0.2          Idle
  10.9.0.3          Idle
```

This table describes the significant fields shown in the display.

Table 40: show protocols (BGP) Field Descriptions

Field	Description
Routing Protocol:	Identifies BGP as the running protocol and displays the BGP AS number. • Range for 2-byte Autonomous system numbers (ASNs) is 1 to 65535. • Range for 4-byte Autonomous system numbers (ASNs) in asplain format is 1 to 4294967295. • Range for 4-byte Autonomous system numbers (ASNs) is asdot format is 1.0 to 65535.65535.
Address Family	Specifies the address family. This can be IPv4 Unicast, IPv4 Multicast, or IPv6 Unicast.
Distance: external	Specifies the distance BGP sets when installing eBGP routes into the RIB. eBGP routes are routes received from eBGP peers. The RIB uses the distance as a tiebreaker when several protocols install a route for the same prefix.
Distance: internal	Specifies the distance BGP sets for routes received from iBGP peers.
Distance: local	Specifies the distance BGP sets for locally generated aggregates and backdoor routes.
Sourced Networks	List of locally sourced networks. These are networks sourced using the **network** command.
Routing information Sources	List of configured BGP neighbors.
Neighbor	Address of a BGP neighbor.
State/Last update received	State of each neighbor and the time since the last update was received from the neighbor if it is established.

show svd role

To display selective VRF download (SVD) role information, use the **show svd role** command in XR EXEC mode.

show svd role

Syntax Description

This command has no keywords or arguments.

Command Default

None.

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

The **show svd role** command output displays name of the line card and role for each address-family in a table.

Task ID

Task ID	Operation
ip-services	read

This example displays the different nodes in a line card and corresponding IPv4, and IPv6 SVD role information:

```
RP/0/RP0/CPU0:router#show svd role
Thu Mar 10 10:45:17.886 PST
Node Name       IPv4 Role            IPv6 Role
-----------------------------------------------------
0/1/CPU0        Core Facing          Not Interested
0/2/CPU0        Core Facing          Core Facing
0/4/CPU0        Standard             Standard
0/5/CPU0        Standard             Standard
```

show svd state

To display selective VRF download (SVD) state information, use the **show svd state** command in XR EXEC mode.

show svd state

Syntax Description This command has no keywords or arguments.

Command Default None.

Command Modes XR EXEC mode

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines No specific guidelines impact the use of this command.

Task ID

Task ID	Operation
ip-services	read

This example shows the SVD configuration state and the SVD operational state in a line card:

```
RP/0/RP0/CPU0:router#show svd state
Thu Mar 10 10:45:32.184 PST
Selective VRF Download (SVD) Feature State:
SVD Configuration State        Enabled
SVD Operational State          Enabled
```

shutdown (BGP)

To disable a neighbor without removing its configuration, use the **shutdown** command in an appropriate configuration mode. To re-enable the neighbor and reestablish a Border Gateway Protocol (BGP) session, use the **no** form of this command.

shutdown [**inheritance-disable**]
no shutdown [**inheritance-disable**]

Syntax Description

inheritance-disable	(Optional) Overrides the value of a **shutdown** command inherited from a neighbor group or session group.

Command Default

Neighbors are not shutdown.

Command Modes

Neighbor configuration

VRF neighbor configuration

Neighbor group configuration

Session group configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

Use the **shutdown** command to terminate any active session for the specified neighbor and remove all associated routing information. Use of the **shutdown** command with a neighbor group or session group may suddenly terminate a large number of BGP neighbor sessions because all neighbors using the neighbor group or session group may be affected.

Use the **show bgp summary** command to display a summary of BGP neighbors. Neighbors that are idle due to the **shutdown** command are displayed with the "Idle (Admin)" state.

If this command is configured for a neighbor group or session group, all neighbors using the group inherit the configuration. Values of commands configured specifically for a neighbor override inherited values.

Task ID

Task ID	Operations
bgp	read, write

Examples

The following example shows that any active session for neighbor 192.168.40.24 is disabled:

```
RP/0/RP0/CPU0:router(config)# router bgp 1
RP/0/RP0/CPU0:router(config-bgp)# neighbor 192.168.40.24
RP/0/RP0/CPU0:router(config-bgp-nbr)# shutdown
RP/0/RP0/CPU0:router(config-bgp-nbr)# exit
```

In the following example, the session remains active for neighbor 192.168.40.24 because the inherited **shutdown** command has been overridden:

```
RP/0/RP0/CPU0:router(config)# router bgp 1
RP/0/RP0/CPU0:router(config-bgp)# session-group group1
RP/0/RP0/CPU0:router(config-bgp-sngrp)# shutdown
RP/0/RP0/CPU0:router(config-bgp-sngrp)# exit
RP/0/RP0/CPU0:router(config-bgp)# neighbor 192.168.40.24
RP/0/RP0/CPU0:router(config-bgp-nbr)# remote-as 1
RP/0/RP0/CPU0:router(config-bgp-nbr)# use session-group group1
RP/0/RP0/CPU0:router(config-bgp-nbr)# shutdown inheritance-disable
RP/0/RP0/CPU0:router(config-bgp-nbr)# exit
```

shutdown (rpki-server)

To shutdown RPKI cache-server, use the **shutdown** command in rpki-server configuration mode. To set that the RPKI cache be active, use the **no** form of this command.

shutdown
no shutdown

This command has no keywords or arguments.

Command Default

RPKI cache is active.

Command Modes

RPKI server configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

Task ID

Task ID	Operation
bgp	read, write

This command shows how to configure no shutdown of the RPKi cache configuration after other RPKI cache parameters are configured:

```
RP/0/RP0/CPU0:router#configure
RP/0/RP0/CPU0:router(config)#router bgp 100
RP/0/RP0/CPU0:router(config-bgp)#rpki server 172.168.35.40
RP/0/RP0/CPU0:router(config-bgp-rpki-cache)# transport ssh port 22
RP/0/RP0/CPU0:router(config-bgp-rpki-cache)#username rpki-user
RP/0/RP0/CPU0:router(config-bgp-rpki-cache)#password rpki-ssh-pass
RP/0/RP0/CPU0:router(config-bgp-rpki-cache)#preference 1
RP/0/RP0/CPU0:router(config-bgp-rpki-cache)#purge-time 30
RP/0/RP0/CPU0:router(config-bgp-rpki-cache)#refresh-time 30
RP/0/RP0/CPU0:router(config-bgp-rpki-cache)#response-time 30
RP/0/RP0/CPU0:router(config-bgp-rpki-cache)#no shutdown
```

signalling disable

To disable BGP or LDP signaling protocol to neighbors, use the **signalling disable** command in neighbor address family (l2vpn vpls-vpws) configuration mode. To restore the system to its default condition, use the **no** form of this command.

signalling {**bgp** | **ldp**} **disable**

Syntax Description

bgp	Selects BGP signaling protocol to disable.
ldp	Selects LDP signaling protocol to disable.

Command Default

Both BGP and LDP signaling protocols are enabled.

Command Modes

Neighbor address family configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

Task ID

Task ID	Operation
bgp	read, write

Example

The following example shows how to disable BGP signaling protocol for neighbor 10.2.3.4:

```
RP/0/RP0/CPU0:router(config)#router bgp 100
RP/0/RP0/CPU0:router(config-bgp)#neighbor 10.2.3.4
RP/0/RP0/CPU0:router(config-bgp-nbr)#address-family l2vpn vpls-vpws
RP/0/RP0/CPU0:router(config-bgp-nbr-af)#signalling bgp disable
```

site-of-origin (BGP)

To attach a site-of-origin extended community attribute to each route received from the specified peer, use the **site-of-origin** command in VRF neighbor address family configuration mode. To restore the system to its default condition, use the **no** form of this command.

site-of-origin [*as-number:nn ip-address:nn*]

Syntax Description

as-number:nn	• *as-number*— Autonomous system (AS) number. • Range for 2-byte Autonomous system number is 1 to 65535. • Range for 4-byte Autonomous system number in asplain format is 1 to 4294967295. • Range for 4-byte Autonomous system number is asdot format is 1.0 to 65535.6553. • *nn* —32-bit number
ip-address:nn	IP address. • *ip-address* —32-bit IP address • *nn* —16-bit number

Command Default

No default behavior or values

Command Modes

VRF neighbor address family configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

When routes are advertised to the peer, routes whose extended communities list contain the site of origin (SoO) are filtered out and not advertised to the peer. Site-of-origin uniquely identifies the site from which the provide edge (PE) router learned routes, thus filtering based on the extended community helps prevent transient routing loops from occurring in complex and mixed network topologies.

Task ID

Task ID	Operations
bgp	read, write

Examples

The following example shows how to configure SoO filtering:

```
RP/0/RP0/CPU0:router(config)# router bgp 6
RP/0/RP0/CPU0:router(config-bgp)# vrf vrf_A
RP/0/RP0/CPU0:router(config-bgp-vrf)# neighbor 192.168.70.24
```

```
RP/0/RP0/CPU0:router(config-bgp-vrf-nbr)# remote-as 10
RP/0/RP0/CPU0:router(config-bgp-vrf-nbr)# address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-bgp-vrf-nbr-af)# site-of-origin 10.0.01:20
```

slow peer (BGP router global configuration)

To enable BGP global slow peer configuration, use the **slow-peer** command in BGP router configuration mode. Enabling global slow peer configuration enables slow peer configuration for all (default VRF and non-default VRF) BGP neighbor address-families.

The default global slow peer mode is *detection-only*.

Use the **no** form of the command to return to the default slow peer configuration.

slow-peer {dynamic [threshold *sec*] | detection-disable}

Syntax Description

dynamic	Enables dynamic slow peer configuration on all (default VRF and non-default VRF) BGP neighbor address-families.
threshold	(Optional) Threshold value (in seconds) used in detecting slow peer. Range is 120 to 3600. Default is 300.
detection-disable	Disables detection-only slow peer configuration on all (default VRF and non-default VRF) BGP neighbor address-families.

Command Default

detection-only

threshold: 300

Command Modes

BGP Router configuration mode

Command History

Release	Modification
Release 7.9.1	This command was introduced.

Usage Guidelines

You can enable slow peer configuration either by using global router configuration mode or by using neighbor address-families.

This section describes global slow peer configuration under router configuration. See below for the effective slow peer configuration when slow peer is configured under both global and neighbor address-families.

-		Global configuration		
-		**[None]**	**[Dynamic]**	**[Detection disable]**
Neighbor address-family configuration	**[None]**	*Detection-only*	*Dynamic*	*None*
	[Static]	*Static*	*Static*	*Static*
	[Dynamic]	*Dynamic*	*Dynamic*	*Dynamic*
	[Dynamic Disable]	*Detection-only*	*None*	*None*

Slow peer global configuration can be configured as *dynamic* or *detection-disable*. The default is *detection-only*.

When global slow peer configuration is not enabled, the default slow peer functionality is *detection-only*. It means all (default VRF and non-default VRF) BGP neighbor address-families are operating in *detection-only* mode. In *detection-only* mode of operation, whenever a neighbor address-family is detected as slow or recovers from being slow, an IOS message is displayed, but there will not be any mitigation to handle slow peers.

When slow peer global configuration is *detection-disable*, then slow peer processing is disabled on all (default VRF and non-default VRF) BGP neighbor address-families.

When the slow peer global configuration is *dynamic*, all VRF BGP neighbor address-families, default and non-default, operate as dynamic slow peers. In the *dynamic* slow peer mode of operation, whenever a neighbor address-family is detected as being slow, in addition to displaying an IOS message, the peer is moved to its own refresh sub-group without impacting other peers in the update-group or a sub-group.

Task ID

Task ID	Operation
bgp	read, write

This example below shows how to enable dynamic slow peer on all (default VRF and non-default VRF) BGP neighbor address-families:

```
Router#configure
Router(config)#router bgp 100
Router(config-bgp)#slow-peer dynamic
Router(config-bgp)#commit
```

This example below shows how to disable slow peer on all (default VRF and non-default VRF) BGP neighbor address-families:

```
Router#configure
Router(config)#router bgp 100
Router(config-bgp)#slow-peer detection-disable
Router(config-bgp)#commit
```

This example below shows how to enable dynamic slow peer with detection threshold of 120 seconds on all (default VRF and non-default VRF) BGP neighbor address-families:

```
Router#configure
Router(config)#router bgp 100
Router(config-bgp)#slow-peer dynamic threshold 120
Router(config-bgp)#commit
```

slow peer (BGP neighbor address-family configuration)

To enable slow peer for a BGP neighbor address-family, use **slow-peer** command in neighbor address-family configuration mode.

By default, BGP neighbor address-family slow peer configuration is disabled.

Use the **no** form of the command to return to the default slow peer configuration.

slow-peer {static | dynamic [threshold *seconds*] | dynamic disable}

Syntax Description

static	Enables static slow peer configuration of a BGP neighbor address-family.
dynamic	Enables dynamic slow peer configuration of a BGP neighbor address-family.
threshold	(Optional) Threshold value (in seconds) used in detecting slow peer. Range is 120 to 3600. Default is 300. The detection logic is evaluated every scanner run whose default interval is 60 seconds.
dynamic disable	Disables dynamic slow peer configuration of a BGP neighbor address-family.

Command Default

threshold: 300

Command Modes

BGP Neighbor address-family configuration

BGP VRF Neighbor address-family configuration

Command History

Release	Modification
Release 7.9.1	This command was introduced.

Usage Guidelines

You can enable slow peer configuration either by using global router configuration mode or by using neighbor address-families.

This section describes slow peer configuration under neighbor address-family configuration. See below for the effective slow peer configuration state when slow peer is configured under both global and neighbor address-family.

-	Global configuration		
-	**[None]**	**[Dynamic]**	**[Detection disable]**

-		Global configuration		
Neighbor address-family configuration	**[None]**	*Detection-only*	*Dynamic*	*None*
	[Static]	*Static*	*Static*	*Static*
	[Dynamic]	*Dynamic*	*Dynamic*	*Dynamic*
	[Dynamic Disable]	*Detection-only*	*None*	*None*

Slow peer configuration under neighbor address-family can be configured as:

- *static*
- *dynamic*
- *dynamic disable*

When slow peer is not configured under neighbor address-family, slow peer handling is disabled for that BGP neighbor address-family.

When slow peer *static* is configured under a neighbor address-family, then that neighbor address-family is moved into its own unique update-group, thus isolating this neighbor address-family from other neighbors. If the user's intention is to group all the slow-peers into a single update group, it can be accomplished by removing static slow peer configuration and configuring the same neighbor out route-policy for all the neighbors.

When slow peer *dynamic* is configured under the neighbor address-family, that BGP neighbor address-family is enabled for dynamic slow peer processing. When the neighbor address-family is enabled for dynamic slow peer processing, whenever the neighbor address-family is detected as slow, the neighbor address-family is processed in its own refresh sub-group without affecting other neighbors in the sub-group, in addition to displaying an IOS message indicating the neighbor address-family has become slow.

When slow-peer *dynamic disable* is configured under the neighbor address-family, it disables the dynamic slow peer processing for that neighbor address-family if dynamic slow peer processing was enabled due to global slow peer *dynamic* configuration.

Task ID

Task ID	Operation
bgp	read, write

This example below shows how to configure static slow peer for a (default VRF and non-default VRF) BGP neighbor address-family:

```
Router#configure
Router(config)#router bgp 100
Router(config-bgp)#neighbor 50.0.0.1
Router(config-bgp-nbr)#address-family ipv4 unicast
Router(config-bgp-nbr-af)#slow-peer static
Router(config-bgp-nbr-af)#commit
```

This example below shows how to disable slow peer for a (default VRF and non-default VRF) BGP neighbor address-family:

```
Router#configure
Router(config)#router bgp 100
Router(config-bgp)#neighbor 50.0.0.1
Router(config-bgp-nbr)#address-family ipv4 unicast
Router(config-bgp-nbr-af)#slow-peer dynamic disable
Router(config-bgp-nbr-af)#commit
```

This example below shows how to enable dynamic slow peer for a (default VRF and non-default VRF) BGP neighbor address-family:

```
Router#configure
Router(config)#router bgp 100
Router(config-bgp)#neighbor 50.0.0.1
Router(config-bgp-nbr)#address-family ipv4 unicast
Router(config-bgp-nbr-af)#slow-peer dynamic
Router(config-bgp-nbr-af)#commit
```

This example below shows how to enable dynamic slow peer with detection threshold of 120 seconds for a (default VRF and non-default VRF) BGP neighbor address-family:

```
Router#configure
Router(config)#router bgp 100
Router(config-bgp)#neighbor 50.0.0.1
Router(config-bgp-nbr)#address-family ipv4 unicast
Router(config-bgp-nbr-af)#slow-peer dynamic threshold 120
Router(config-bgp-nbr-af)#commit
```

socket receive-buffer-size

To set the size of the receive buffers for all Border Gateway Protocol (BGP) neighbors, use the **socket receive-buffer-size** command in an appropriate configuration mode. To set the size of the receive buffers to the default size, use the **no** form of this command.

socket receive-buffer-size *socket-size* [*bgp-size*]
no socket receive-buffer-size [*socket-size*] [*bgp-size*]

Syntax Description

socket-size	Size (in bytes) of the receive-side socket buffers. Range is 512 to 131072.
bgp-size	(Optional) Size (in bytes) of the receive buffers in BGP. Range is 512 to 131072.

Command Default

socket-size : 32,768 bytes

bgp-size : 4,032 bytes

Command Modes

Router configuration

VRF configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

Use the **socket receive-buffer-size** command to increase the buffer size when receiving updates from a neighbor. Using larger buffers can improve convergence time because the software can process more packets simultaneously. However, allocating larger buffers uses more memory on your router.

Note Increasing the socket buffer size uses more memory only when more messages are waiting to be processed by the software. In contrast, increasing the BGP buffer size uses extra memory indefinitely.

Use the **receive-buffer-size** command on individual neighbors to change the values set by the **socket receive-buffer-size** command.

Task ID

Task ID	Operations
bgp	read, write

Examples

The following example shows how to set the receive buffer sizes for all neighbors to 65,536 bytes for the socket buffer and 8192 bytes for the BGP buffer:

```
RP/0/RP0/CPU0:router(config)# router bgp 1
```

```
RP/0/RP0/CPU0:router(config-bgp)# socket receive-buffer-size 65536 8192
```

socket send-buffer-size

To set the size of the send buffers for all Border Gateway Protocol (BGP) neighbors, use the **socket send-buffer-size** command in an appropriate configuration mode. To set the size of the send buffers to the default size, use the **no** form of this command.

socket send-buffer-size *socket-size* [*bgp-size*]
no socket send-buffer-size [*socket-size*] [*bgp-size*]

Syntax Description

socket-size	Size (in bytes) of the send-side socket buffers. Range is 4096 to 131072.
bgp-size	(Optional) Size (in bytes) of the send buffers in BGP. Range is 4096 to 131072.

Command Default

socket-size : 10240 bytes

bgp-size : 4096 bytes

Command Modes

Router configuration

VRF configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

Use the **socket send-buffer-size** command to increase the buffer size when sending updates to neighbors. Using larger buffers can improve convergence time because the software can process more packets simultaneously. However, allocating larger buffers uses more memory on your router.

Note Increasing the socket buffer size uses more memory only when more messages are waiting to be sent by the software. In contrast, increasing the BGP buffer size uses extra memory indefinitely.

Use the **send-buffer-size** command on individual neighbors to change the values set by the **socket send-buffer-size** command.

Task ID

Task ID	Operations
bgp	read, write

Examples

The following example shows how to set the send buffer sizes for all neighbors to 8192 bytes for the socket buffer and the BGP buffer:

```
RP/0/RP0/CPU0:router(config)# router bgp 1
```

```
RP/0/RP0/CPU0:router(config-bgp)# socket send-buffer-size 8192 8192
```

soft-reconfiguration inbound

To configure the software to store updates received from a neighbor, use the **soft-reconfiguration inbound** command in an appropriate configuration mode. To disable storing received updates, use the **no** form of this command.

soft-reconfiguration inbound [**always** | **inheritance-disable** | **RPKI-dropped-only** | **RPKI-tested-only**]
no soft-reconfiguration inbound [**always** | **inheritance-disable** | **RPKI-dropped-only** | **RPKI-tested-only**]

Syntax Description

always	(Optional) Saves a copy of all received routes, even if the neighbor supports the route refresh capability.
inheritance-disable	(Optional) Overrides configuration for this command that may be inherited from a neighbor group or address family group.
RPKI-dropped-only	(Optional) Saves a copy of only the routes dropped by an RPKI validation-state test in neighbor-in route-policy.
RPKI-tested-only	(Optional) Saves a copy of only the routes tested in an RPKI validation-state test in neighbor-in route-policy.

Command Default

Soft reconfiguration is not enabled.

Command Modes

IPv4 address family group configuration

IPv6 address family group configuration

IPv4 neighbor address family configuration

IPv4 neighbor group address family configuration

IPv6 neighbor group address family configuration

VPNv4 neighbor address family configuration

VPNv4 address family group configuration

VRF IPv4 neighbor address family configuration

VPNv4 neighbor group address family configuration

VPNv6 address family group configuration

VPNv6 neighbor address family configuration

VRF IPv6 neighbor address family configuration

VPNv6 neighbor group address family configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

To filter or modify some of the updates received from a neighbor, you configure an inbound policy using the **route-policy (BGP)** command. Configuring soft reconfiguration inbound causes the software to store the original unmodified route beside a route that is modified or filtered. This allows a "soft clear" to be performed after the inbound policy is changed. To perform a soft clear, use the **clear bgp soft** command with the **in** keyword specified. The unmodified routes are then passed through the new policy and installed in the BGP table.



Note If an address family group, neighbor group, or session group is configured, the configuration inside these configuration groups will not be effective unless it is applied directly or indirectly to one or more neighbors.



Note The bgp auto-policy-soft-reset is enabled by default. A soft clear is done automatically when the inbound policy configured with the **route-policy** (BGP) command is changed. This behavior can be changed by disabling the auto-policy-soft-reset using the **bgp auto-policy-soft-reset disable** command.

If the neighbor supports the route refresh capability, then the original routes are not stored because they can be retrieved from the neighbor through a route refresh request. However, if the **always** keyword is specified, the original routes are stored even when the neighbor supports the route refresh capability.

If the **soft-reconfiguration inbound** command is not configured and the neighbor does not support the route refresh capability, then an inbound soft clear is not possible. In that case, the only way to rerun the inbound policy is to use the **clear bgp** *ip-address* command to reset the neighbor BGP session.



Note If there is an existing BGP session with a neighbor that does not support the route refresh capability, the session is terminated and a new one is initiated.



Note The extra routes stored as a result of configuring this command use more memory on the router.

If you configure this command for a neighbor group or neighbor address family group, all neighbors using the group inherit the configuration. Values of commands configured specifically for a neighbor override inherited values.

Task ID

Task ID	Operations
bgp	read, write

Examples

The following example shows inbound soft reconfiguration enabled for IP Version 4 (IPv4) unicast routes received from neighbor 10.108.1.1. The software stores all routes received in their unmodified form so that when an inbound soft clear is performed later, the stored information can then be used to generate a new set of modified routes.

```
RP/0/RP0/CPU0:router(config)# router bgp 100
RP/0/RP0/CPU0:router(config-bgp)# neighbor 10.108.1.1
RP/0/RP0/CPU0:router(config-bgp-nbr)# remote-as 100
RP/0/RP0/CPU0:router(config-bgp-nbr)# address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-bgp-nbr-af)# soft-reconfiguration inbound
RP/0/RP0/CPU0:router(config-bgp-nbr-af)# exit
```

speaker-id

To allocate a speaker process to a neighbor, use the **speaker-id** command in the appropriate configuration mode. To remove the speaker process from a neighbor, use the **no** form of this command.

speaker-id *id*
no speaker-id [*id*]

Syntax Description

id	ID of the speaker process. Range is 1 to 15.

Command Default

Default is 0.

Command Modes

Neighbor configuration

Session group configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

Task ID

Task ID	Operations
bgp	read, write

Examples

The following example shows how to allocate speaker process 3 to neighbor 192.168.40.24:

```
RP/0/RP0/CPU0:router(config)# router bgp 109
RP/0/RP0/CPU0:router(config-bgp)# neighbor 192.168.40.24
RP/0/RP0/CPU0:router(config-bgp-nbr)# speaker-id 3
```

stats resource reassign

To reassign unused counter engines of one feature with counter engine of other features, use the **stats-resource reassign** command in XR Config mode. Before reassigning the NPU engine counters, you must make sure that they are unused. You could use the **show controllers npu resources stats internal instance** *npu-id* **location** *linecard-id* command to view the statistics for the NPU resources.

stats-resource reassign location *linecard-id* **internal feature** [**aclrx-mcrouterx** | **acltx** | **l2rx** | **l2tx-mctx** | **l3rx** | **l3tx** | **lpts-pbr** | **mplsrx** | **policer** | **tetx** | **trap** | **voq**] **eng-ids** *ID*

Syntax Description

location *linecard-id*	Specify the linecard-id of the counter engine to be reassigned.
internal feature	Specifies the internal feature to be configured. **Note** Internal feature is currently supported for internal counter only.
aclrx-mcrouterx	Specifies the Ingress ACL Rx and McRoute feature option
acltx	Specifies the Egress ACL Tx feature option
l2rx	Specifies the Ingress L2 Rx Interface feature option
l2tx-mctx	Specifies the Egress L2 Tx, BVI, and McTx Interface feature option
l3rx	Specifies the Ingress L3 Rx and BVI Interface feature option
l3tx	Specifies the Egress L3 Tx Interface feature option
lpts-pbr	Specifies the LPTS and PBR feature option
mplsrx	Specifies the SR Rx feature option
policer	Specifies the Ingress QoS feature option
tetx	Specifies the MPLS Te Tx feature option
trap	Specifies the trap feature option
voq	Specifies the VoQ feature option
eng-ids *ID*	Specify the engine IDs to be reassigned. You could reassign one or more engines simultaneously by mentioning the set of IDs separated by space.

Command Default

The NPU counters engines are assigned by default based on the feature scales supported.

Command Modes

XR Config mode

Command History

Release	Modification
Release 7.6.1	This command was introduced.
Release 7.6.2	The **location** keyword was introduced.

Usage Guidelines

- This command is supported only on NC57-24DD-** and NC57-18DD-SE-** line cards.
- NC57 line cards support 22 counter engines with engine identifiers ranging from 0 - 21. You can reassign one or more engines simultaneously by mentioning the set of IDs separated by space.
- Re-assignment of counter engines is not supported for Policer. The Policer feature includes QoS Rx and Policer metering.
- When EnhQoS is enabled, the counter set size for QoS Rx is four and when Wred is enabled, the counter set size for VOQ is six.
- The aclrx-mcrouterx feature is associated with ACL access-group stats and the hw-module route-stats l3mcast IPv4 command

Note The stats resource reassign command is a replacement for the deprecated 7.6.1 command hw-module profile stats j2-dynamic-stats.

Table 41: Counter Engine Capacity

Engine ID	Capacity
0 - 7	4K
8 - 15	8K
16 - 21	16K

Note The counter engine allocation differs on each platform and profile.

The following table shows a sample counter engine allocation for L3MAX-Non SE:

Table 42: L3MAX-Non SE default feature engine matrix

Feature	Scale	Default counter set	Default counter engine
trap	• TRAP = 2K	2	• 4K - Engine ID 0
lpts-pbr	• LPTS = 4K • PRB = 4K	2	• 8K - Engine ID 8

Feature	Scale	Default counter set	Default counter engine
mplsrx	• SR Rx = 4K	1	• 4K - Engine ID 1
aclrx-mcrouterx	• ACLRx-McRoute = 8K	1	• 8K - Engine ID 9
tetx	• TeTx = 8K	1	• 8K - Engine ID 10
acltx	• ACLTx = 4K	1	• 4K - Engine ID 2
l2rx	• L2Rx = 8K	1	• 8K - Engine ID 11
l2tx-mctx	• L2Tx+BVI-Tx = 16K • McTx (V4+V6) = 4K	1	• 4K - Engine ID 3 • 8K - Engine ID 12 • 8K - Engine ID 13
l3rx	• L3Rx = 2K • BVI-Rx = 2.8K	1	• 4K - Engine ID 4 • 4K - Engine ID 5 • 8K - Engine ID 14 • 8K - Engine ID 15
l3tx	• L3Tx = 4K	1	• 4K - Engine ID 6
l3tx_2	• L3Tx_2 = 2K	2	• 4K - Engine ID 7
voq	• 2K interface with class map size 8 = 16K	2	• 16K - Engine ID 16 • 16K - Engine ID 17
policer	• 2K interface with class map size 8 = 16K	2	• 16K - Engine ID 18 • 16K - Engine ID 19
Metering	• Parent = 16K • Child = 16K	1	• 16K - Child - Engine ID 20 • 16K - Parent - Engine ID 21

The following example shows how to verify the assigned and unused engines:

```
Router(config)# sh controllers npu resources stats internal instance all location 0/0/CPU0

Fri Sep 16 11:43:16.543 UTC
```

```
System information for NPU 0:
Counter Processor Configuration Profile: Default
================================================================================================
Assigned Counter Engine Resource Information
================================================================================================
Feature                  Eng    State    Eng   Total Counter    Core-0 Entries  Core-1
Entries
                          Id               Size    Entries          InUse           InUse

================================================================================================
trap-aclpolicer          0      In use   4K     2048            143              143

lpts-pbr                 1      In use   4K     2048            6                6

lpts-pbr                 2      Free     4K     2048            0                0

mplsrx                   3      Free     4K     4096            0                0

mplsrx                   4      Free     4K     4096            0                0

mplsrx                   5      Free     4K     4096            0                0

l2rx                     6      Free     4K     4096            0                0

acltx                    7      Free     4K     4096            0                0

lpts-pbr                 8      Free     8K     4096            0                0

lpts-pbr                 9      Free     8K     4096            0                0

aclrx-mcrouterx          10     Free     8K     8192            0                0

aclrx-mcrouterx          11     Free     8K     8192            0                0

aclrx-mcrouterx          12     Free     8K     8192            0                0

tetx                     13     Free     8K     8192            0                0

tetx                     14     Free     8K     8192            0                0

l2rx                     15     Free     8K     8192            0                0

POLICER1 (CHILD)              16    Reserved 16K    16384          0               0

POLICER1 (CHILD)              17    Reserved 16K    16384          0               0

POLICER1 (CHILD)              18    Reserved 16K    16384          0               0

POLICER1 (CHILD)              19    Reserved 16K    16384          0               0

Storm Control Policer         20    Reserved 16K    16384          0               0

POLICER2 (PARENT)             21    Reserved 16K    16384          0               0

================================================================================
Note: Free and Unassigned Counter engines can be reassigned to other features
================================================================================
```

Examples

The following example shows how to add counter engines to a features:

```
Router(config)# stats-resource reassign location 0/0/CPU0
Router(config)# internal feature lpts-pbr eng-ids 6 16 17
Router(config)# commit
```

The following example shows how to replace old engines with new engines for a feature.

The counter engine 17 for lpts-pbr feature is replaced with engine 18.

```
Router(config)# stats-resource reassign location 0/0/CPU0
Router(config)# internal feature lpts-pbr eng-ids 6 16 18
Router(config)# commit
```

The following example shows how to scale down engines for a feature.

The counter engine 18 for lpts-pbr feature is unassigned and moved to unassigned counter engines.

```
Router(config)# stats-resource reassign location 0/0/CPU0
Router(config)# internal feature lpts-pbr eng-ids 6 16
Router(config)# commit
```

The following example shows how to unassign all the engines which were configured for the feature:

```
Router(config)# stats-resource reassign location 0/0/CPU0
Router(config)# no internal feature lpts-pbr eng-ids
Router(config)# commit
```

svd platform enable

To enable selective VRF download (SVD) for Cisco ASR 9000 Series Aggregation Services Router, use the **svd platform enable** command in Admin Configuration modeSystem Admin Config mode. To disable selective VRF download for Cisco ASR 9000 Series Aggregation Services Router, use the **no** form of this command.

svd platform enable
no svd platform enable

Syntax Description

This command has no keywords or arguments.

Usage Guidelines

After upgrading to Cisco IOS XR Release 4.3.1 or later, the default setting for SVD is disabled and routes or labels are not selectively downloaded to the line cards. All routes are be available on all line cards. This may lead to out of resource conditions, if the line card hardware scale limits are exceeded because routes were conserved by downloading selectively earlier. To resolve the out of resource condition, turn on SVD using the **svd platform enable** command.

Ensure that the total number of routes (sum of all unique routes on core facing line cards and edge facing line cards) fits on the line card hardware, before upgrading to Cisco IOS XR Release 4.3.1 or later.

To enable selective VRF download on NCS 5500 Series Routers, reload the chassis using the **reload location all** command after configuring the **svd platform enable** command.

Note After enabling SVD using svd platform enable, do not use the **selective-vrf-download disable** command to turn off SVD.

Selective VRF download is disabled by default. Once SVD is enabled, use the **no svd platform enable** command followed by a reload of the router using the **reload location all** command to disable SVD.

Command Default

Selective VRF download is disabled.

Command Modes

Admin Configuration modeSystem Admin Config mode

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

Task ID

Task ID	Operation
config-services	read, write

This example shows how to enable selective VRF download for Cisco ASR 9000 Series Aggregation Services Router:

```
RP/0/RP0/CPU0:router#configure
RP/0/RP0/CPU0:router(config)#svd platform enable
```

table-policy

To apply a routing policy to routes being installed into the routing table, use the **table-policy** command in an appropriate configuration mode. To disable applying a routing policy when installing routes into the routing table, use the **no** form of this command.

table-policy *policy-name*
no table-policy [*policy-name*]

Syntax Description

policy-name	Name of the routing policy to apply.

Command Default

No policy is applied when routes are installed into the routing table.

Command Modes

IPv4 address family configuration

IPv6 address family configuration

VRF IPv4 address family configuration

VRF IPv6 address family configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

Note Table policy provides users with the ability to drop routes from the RIB based on match criteria. This feature can be useful in certain applications and should be used with caution as it can easily create a routing 'black hole' where BGP advertises routes to neighbors that BGP does not install in its global routing table and forwarding table.

Use the **table-policy** command to modify route attributes as the routes are installed into the routing table by Border Gateway Protocol (BGP). Commonly, it is used to set the traffic index attribute.

Task ID

Task ID	Operations
bgp	read, write

Examples

The following example shows how to apply the set-traffic-index policy to IPv4 unicast routes being installed into the routing table:

```
RP/0/RP0/CPU0:router(config)# router bgp 1
RP/0/RP0/CPU0:router(config-bgp)# address-family ipv4 unicast
```

```
RP/0/RP0/CPU0:router(config-bgp-af)# table-policy set-traffic-index
```

timers (BGP)

To set the timers for a specific Border Gateway Protocol (BGP) neighbor, use the **timers** command in an appropriate configuration mode. To set the timers to the default values, use the **no** form of this command.

timers *keepalive hold-time*
no timers [*keepalive hold-time*]

Syntax Description

keepalive	Frequency (in seconds) with which the software sends keepalive messages to a neighbor. Range is 0 to 65535.
hold-time	Interval (in seconds) after not receiving a keepalive message from the neighbor that the software terminates the BGP session for the neighbor. Values are 0 or a number in the range from 3 to 65535.

Command Default

keepalive : 60 seconds

hold-time : 180 seconds

Use the **timers bgp** command to override the default values.

Command Modes

Neighbor configuration

VRF neighbor configuration

Neighbor group configuration

Session group configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

The timers actually used in connection with the neighbor may not be the same as those configured with this command. The actual timers are negotiated with the neighbor when establishing the session. The negotiated hold time is the minimum of the configured time and the hold time received from the neighbor. If the negotiated hold time is 0, keepalives are disabled.

The configured value for the keepalive must not exceed one-third of the negotiated hold time. If it does, a value of one-third of the negotiated hold time is used.

If this command is configured for a neighbor group or neighbor address family group, all neighbors using the group inherit the configuration. Values of commands configured specifically for a neighbor override inherited values.

In cases where mechanisms such as Bi-directional Forwarding Detection (BFD), BGP fast-external-failover or Next-hop Tracking cannot be employed to detect and react to changes in the network in a faster manner, BGP Keepalive and Hold-timer values can be configured to use smaller values than the default (60 and 180 seconds respectively). When using aggressive values, consider the router's profile and scale, particularly in respect to the number of BGP neighbours that will be using sessions with the non-default timers.

Sessions using very aggressive values will be more susceptible to flap during events that cause the Route-Processor's CPU utilization levels to increase. Such events include component OIR, Route-Processor Failover, network instability, excessive churn in routing protocols etc. It is therefore recommended that the desired scale and profile of the router be tested with the non-default timer values, subjecting the router to CPU-intensive events in order to determine the timer threshold values that are appropriate for the router before configuring the values in an operational network.

The BGP Non-Stop Routing (NSR) is able to sustain sessions with more aggressive timer values than BGP Graceful Restart (GR) since in the event of a Route-Processor Failover, Graceful Restart (GR) requires the re-establishment of the TCP session over which the BGP session takes place. When using Non-Stop Routing (NSR), both the underlying TCP session and BGP session are maintained during Route-Processor failover.

Task ID

Task ID	Operations
bgp	read, write

Examples

The following example shows how to change the keepalive timer to 70 seconds and the hold-time timer to 210 seconds for the BGP peer 192.168.40.24:

```
RP/0/RP0/CPU0:router(config)# router bgp 109
RP/0/RP0/CPU0:router(config-bgp)# neighbor 192.168.40.24
RP/0/RP0/CPU0:router(config-bgp-nbr)# remote-as 1
RP/0/RP0/CPU0:router(config-bgp-nbr)# timers 70 210
```

timers bgp

To change the default timer values for Border Gateway Protocol (BGP) neighbors, use the **timers bgp** command in an appropriate configuration mode. To set the default timers to the default values, use the **no** form of this command.

timers bgp *keepalive hold-time*
no timers bgp [*keepalive hold-time*]

Syntax Description

keepalive	Frequency (in seconds) with which the software sends keepalive messages to a neighbor. Range is 0 to 65535.
hold-time	Interval (in seconds) after not receiving a keepalive message from the neighbor that the software terminates the BGP session for the neighbor. Values are 0 or a number in the range from 3 to 65535.

Command Default

keepalive : 60 seconds

hold-time : 180 seconds

Command Modes

Router configuration

VRF configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

Use the **timers bgp** command to adjust the default timer times used by all BGP neighbors. The values can be overridden on particular neighbors using the **timers** command in the neighbor configuration mode.

The timers actually used in connection with the neighbor may not be the same as those configured with this command. The actual timers are negotiated with the neighbor when establishing the session. The negotiated hold time is the minimum of the configured time and the hold time received from the neighbor. If the negotiated hold time is 0, keepalives are disabled.

The configured value for the keepalive must not exceed one-third of the negotiated hold time. If it does, a value of one-third of the negotiated hold time is used.

In cases where mechanisms such as Bi-directional Forwarding Detection (BFD), BGP fast-external-failover or Next-hop Tracking cannot be employed to detect and react to changes in the network in a faster manner, BGP Keepalive and Hold-timer values can be configured to use smaller values than the default (60 and 180 seconds respectively). When using aggressive values, consider the router's profile and scale, particularly in respect to the number of BGP neighbors that will be using sessions with the non-default timers.

Sessions using very aggressive values will be more susceptible to flap during events that cause the Route-Processor's CPU utilization levels to increase. Such events include component OIR, Route-Processor Failover, network instability, excessive churn in routing protocols etc. It is therefore recommended that the desired scale and profile of the router be tested with the non-default timer values, subjecting the router to CPU-intensive events in order to determine the timer threshold values that are appropriate for the router before configuring the values in an operational network.

The BGP Non-Stop Routing (NSR) is able to sustain sessions with more aggressive timer values than BGP Graceful Restart (GR) since in the event of a Route-Processor Failover, Graceful Restart (GR) requires the re-establishment of the TCP session over which the BGP session takes place. When using Non-Stop Routing (NSR), both the underlying TCP session and BGP session are maintained during Route-Processor failover.

Task ID

Task ID	Operations
bgp	read, write

Examples

The following example shows how to configure a default keepalive time of 30 seconds and a default hold time of 90 seconds:

```
RP/0/RP0/CPU0:router(config)# router bgp 1
RP/0/RP0/CPU0:router(config-bgp)# timers bgp 30 90
```

transport (rpki-server)

To choose a transport mechanism for the RPKI cache-server configuration, establish and manage transport connections, and send or receive byte streams from the network, use the **transport** command in rpki-server configuration mode. To disable the transport connection, use the **no** form of this command.

transport {**ssh** | **tcp**} **port** *port-number*
no transport {**ssh** | **tcp**} **port** *port-number*

Syntax Description

port	Specifies to choose a port number for the RPKI cache transport.
port-number	Specifies the port number for the RPKI cache transport over TCP and SSH protocols. The port number ranges from 1 to 65535. **Note** SSH supports custom ports in addition to the default port number 22.

Command Default

Transport mechanism is disabled.

Command Modes

RPKI server configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

The transport can be set to either TCP or SSH. An SSH transport session is the recommended transport between router and RPKI cache for security reasons.

The transport method (TCP or SSH) can be configured on a per-RPKI-server basis. Changing the transport method will cause the cache session to flap (cleanup its existing transport related data and initialize the new transport related data).

Task ID

Task ID	Operation
bgp	read, write

This example shows how to configure SSH as the transport mechanism and to use port 22 for SSH communication:

```
RP/0/RP0/CPU0:router#configure
RP/0/RP0/CPU0:router(config)#router bgp 100
RP/0/RP0/CPU0:router(config-bgp)#rpki server 172.168.35.40
RP/0/RP0/CPU0:router(config-bgp-rpki-cache)# transport ssh port 22
```

update limit

To set upper bound on transient memory usage for update generation, use the **update limit** command in XR Config mode. To return the bounds to the default value, use the **no** form of this command.

update limit *update-limit-MB*
no update limit

Syntax Description

update-limit-MB	Sets the update limit in megabytes (MB). Range is 16 to 2048 MB.

Command Default

Default update limit is 512 MB.

Command Modes

XR Config mode

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

Use the **update limit** command to configure a global limit on the size of messages the software queues when updating peers. Increasing the limit can result in faster Border Gateway Protocol (BGP) convergence, but also may result in higher memory usage during convergence.

Task ID

Task ID	Operation
bgp	read, write

This example shows how to set the update limit as *1024* MB:

```
RP/0/RP0/CPU0:router(config)# router bgp 65000
RP/0/RP0/CPU0:router(config-bgp)#update limit 1024
```

update limit address-family

To set upper bound on transient memory usage for update generation for an address family, use the **update limit address-family** command in an appropriate address-family configuration mode. To return the bounds to the default value, use the **no** form of this command.

update limit address-family *update-limit-MB*
no update limit address-family

Syntax Description

update-limit-MB	Sets the update limit in megabytes (MB). Range is 4 MB to 2048 MB.

Command Default

Default update limit is 256 MB.

Command Modes

IPv4 address family configuration

IPv6 address family configuration

L2VPN address family configuration

VPNv4 address family configuration

VPNv6 address family configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

Use the **update limit address-family** command to configure a global limit on the size of messages the software queues when updating peers. Increasing the limit can result in faster Border Gateway Protocol (BGP) convergence, but also may result in higher memory usage during convergence.

Task ID

Task ID	Operation
bgp	read, write

This example shows how to set the update limit as *512 MB* for address family IPv4 unicast:

```
RP/0/RP0/CPU0:router#configure
RP/0/RP0/CPU0:router(config)#router bgp 100
RP/0/RP0/CPU0:router(config-bgp)#address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-bgp-af)#update limit address-family 512
```

update limit sub-group

To set upper bound on transient memory usage for update generation for eBGP or iBGP sub-groups, use the **update limit sub-group** command in an appropriate address-family configuration mode. To return the bounds to the default value, use the **no** form of this command.

update limit sub-group {**ebgp** | **ibgp**} *update-limit-MB*
no update limit sub-group {**ebgp** | **ibgp**}

Syntax Description

ebgp	Specifies the update limit for eBGP sub-groups.
ibgp	Specifies the update limit for iBGP sub-groups.
update-limit-MB	Sets the update limit in megabytes (MB). Range is 1 MB to 512 MB.

Command Default

Default update limit is 32 MB.

Command Modes

IPv4 address family configuration

IPv6 address family configuration

L2VPN address family configuration

VPNv4 address family configuration

VPNv6 address family configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

Use the **update limit sub-group** command to configure a global limit on the size of messages the software queues when updating peers. Increasing the limit can result in faster Border Gateway Protocol (BGP) convergence, but also may result in higher memory usage during convergence.

Task ID

Task ID	Operation
bgp	read, write

This example shows how to set the update limit as *256 MB* for eBGP sub-group under address family IPV4 unicast:

```
RP/0/RP0/CPU0:router#configure
RP/0/RP0/CPU0:router(config)#router bgp 100
RP/0/RP0/CPU0:router(config-bgp)#address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-bgp-af)#update limit sub-group ebgp 256
```

update in error-handling basic disable

To disable inbound update message basic error handling for eBGP or iBGP neighbors, use the **update in error-handling basis disable** command in XR Config mode. To enable inbound update message basic error handling, use the **no** form of this command.

update in error-handling basic {**ebgp** | **ibgp**} **disable**
no update in error-handling basic {**ebgp** | **ibgp**} **disable**

Syntax Description

ebgp	Specifies inbound update message basic error handling for eBGP neighbors.
ibgp	Specifies inbound update message basic error handling for iBGP neighbors.

Command Default

Inbound update message basic error handling is enabled.

Command Modes

XR Config mode

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

Task ID

Task ID	Operation
bgp	read, write

This example shows how to disable inbound update message basic error handling for eBGP neighbors:

```
RP/0/RP0/CPU0:router#configure
RP/0/RP0/CPU0:router(config)#router bgp 100
RP/0/RP0/CPU0:router(config-bgp)#update in error-handling basic ebgp disable
```

This example shows how to disable inbound update message basic error handling for iBGP neighbors:

```
RP/0/RP0/CPU0:router#configure
RP/0/RP0/CPU0:router(config)#router bgp 100
RP/0/RP0/CPU0:router(config-bgp)#update in error-handling basic ibgp disable
```

update in error-handling extended

To enable inbound update message extended error handling for eBGP or iBGP neighbors, use the **update in error-handling extended** command in XR Config mode. To disable inbound update message error handling, use the **no** form of this command.

update in error-handling extended {**ebgp** | **ibgp**}
no update in error-handling extended {**ebgp** | **ibgp**}

Syntax Description

ebgp	Specifies to enable inbound update message extended error handling for eBGP neighbors.
ibgp	specifies to enable inbound update message extended error handling for iBGP neighbors.

Command Default

Inbound update message extended error handling is disabled.

Command Modes

XR Config mode

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

Task ID

Task ID	Operation
bgp	read, write

This example shows how to enable inbound update message extended error handling for eBGP neighbors:

```
RP/0/RP0/CPU0:router#configure
RP/0/RP0/CPU0:router(config)#router bgp 100
RP/0/RP0/CPU0:router(config-bgp)#update in error-handling extended ebgp
```

This example shows how to enable inbound update message extended error handling for iBGP neighbors:

```
RP/0/RP0/CPU0:router#configure
RP/0/RP0/CPU0:router(config)#router bgp 100
RP/0/RP0/CPU0:router(config-bgp)#update in error-handling extended ibgp
```

update in error-handling treat-as-withdraw (BGP)

To discard the incoming update message as a withdraw message and to avoid BGP session reset, use the **update in error-handling treat-as-withdraw** command in XR Config mode. To disable withdraw during inbound update message error handling, use the **no** form of this command.

update in error-handling treat-as-withdraw

Syntax Description

This command has no keywords or arguments.

Command Modes

XR Config mode

Command History

Release	Modification
Release 7.10.1	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

Task ID

Task ID	Operation
bgp	read, write

This example shows how to discard the incoming update message as a withdraw message and to avoid BGP session reset:

```
Router(config)#router bgp 109
Router(config-bgp)#neighbor 192.168.40.24
Router(config-bgp-nbr)#update in error-handling treat-as-withdraw
Router(config-bgp-nbr)#commit
```

update in labeled-unicast equivalent

To treat SAFIs unicast and labelled unicast for a given neighbor as same, use the **update in labeled-unicast equivalent** command in neighbor configuration mode. This command was introduced for interoperability between Cisco IOS XR and IOS XE systems. For example, If the XE router advertises ipv4 unicast and then withdraw with the ipv4 label-unicast of the same route, with knob enabled, then the BGP in the receiving XR router treats both of them same. Hence BGP removes the prefix information from its database.

update in labeled-unicast equivalent [inheritance-disable]

Syntax Description

equivalent	IPv4/IPv6 labeled-unicast inbound updates (paths) treated equivalent to unicast updates (paths)
inheritance-disable	Does not inherit this unicast configuration from parent group

Command Default

None

Command Modes

Neighbor configuration mode

Neighbor group configuration mode

Command History

Release	Modification
Release 7.0.1	This command was introduced.

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Task ID

Task ID	Operation
bgp	read, write

Example

The following example shows how to treat IPv4 labeled-unicast inbound updates equivalent to unicast updates (paths):

```
RP/0/RP0/CPU0:router#configure
RP/0/RP0/CPU0:router(config)#router bgp 100
RP/0/RP0/CPU0:router(config-bgp)#neighbor 10.2.3.4
RP/0/RP0/CPU0:router(config-bgp-nbr)#update in labeled-unicast equivalentexit
```

update out logging

To enable logging of update generation events, use the **update out logging** command in XR Config mode. To disable the logging of update generation events, use the **no** form of this command.

update out logging
no update out logging

Syntax Description

This command has no keywords or arguments.

Command Default

Update generation event logging is disabled.

Command Modes

XR Config mode

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

Task ID

Task ID	Operation
bgp	read, write

This example shows how to enable logging of update generation events:

```
RP/0/RP0/CPU0:router#configure
RP/0/RP0/CPU0:router(config)#router bgp 100
RP/0/RP0/CPU0:router(config-bgp)#update out logging
```

update-source

To allow internal Border Gateway Protocol (iBGP) sessions to use the primary IP address from a particular interface as the local address when forming an iBGP session with a neighbor, use the **update-source** command in an appropriate configuration mode. To set the chosen local IP address to the nearest interface to the neighbor, use the **no** form of this command.

update-source *type interface-path-id*
no update-source [*type interface-path-id*]

Syntax Description

type	Interface type. For more information, use the question mark (**?**) online help function.
interface-path-id	Physical interface or virtual interface. **Note** Use the **show interfaces** command to see a list of all interfaces currently configured on the router. For more information about the syntax for the router, use the question mark (**?**) online help function.

Command Default

Best local address

Command Modes

Neighbor configuration

VRF neighbor configuration

Neighbor group configuration

Session group configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

The **update-source** command is commonly used with the loopback interface feature for iBGP sessions. The loopback interface is defined, and the interface address is used as the endpoint for a BGP session through the **update-source** command. This mechanism allows a BGP session to remain up even if the outbound interface goes down, provided there is another route to the neighbor.

If this command is configured for a neighbor group or session group, all neighbors using the group inherit the configuration. Values of commands configured specifically for a neighbor override inherited values.

Task ID

Task ID	Operations
bgp	read, write

Examples

The following example shows how to configure this router to use the IP address from the Loopback0 interface when trying to open a session with neighbor 172.20.16.6:

```
RP/0/RP0/CPU0:router(config)# router bgp 110
RP/0/RP0/CPU0:router(config-bgp)# neighbor 172.20.16.6
RP/0/RP0/CPU0:router(config-bgp-nbr)# remote-as 110
RP/0/RP0/CPU0:router(config-bgp-nbr)# update-source Loopback0
```

update wait-install

To configure BGP to wait for feedback from RIB indicating that the routes that BGP installed in RIB have been installed in FIB, before BGP send out updates to neighbors, use the **update wait-install** command in an appropriate configuration mode.

update wait-install
no update wait-install

This command has no keywords or arguments.

Command Default

The update wait-install configuration is disabled.

Command Modes

Router IPv4 address family

Router VPNv4 address family

Router IPv6 address family

Router VPNv6 address family

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

Starting from Release 6.6.2, when you configure the **wait-install** command, the BGP status message "LC/FIB for the neighbor in reloading" is displayed only when a connection is established with a directly connected EBGP neighbour.

Task ID

Task ID	Operation
bgp	read, write

This example shows how to enable the **update wait-install** configuration under VPNv4 unicast address family:

```
RP/0/RP0/CPU0:router#configure
RP/0/RP0/CPU0:router(config)#router bgp 65500
RP/0/RP0/CPU0:router(config-bgp)#address-family vpnv4 unicast
RP/0/RP0/CPU0:router(config-bgp-af)#update wait-install
```

update wait-install delay startup

To schedule the delay for BGP update generations to prevent routes from being advertised to peers until RIB is synchronized, run the **update wait-install delay startup** command in the System Admin Config mode.

Use the **no** form of this command to disable the feature.

update wait-install delay startup *seconds*

Syntax Description

seconds	Delay for the BGP start-up phase. The range of the delay is from 1 second to 600 seconds.

Command Default

The command is disabled by default.

Command Modes

System Admin Config mode

Command History

Release	Modification
Release 7.5.3	This command was introduced.

Usage Guidelines

This command is applicable for the following Address Family Indicators (AFIs):

- IPv4 unicast
- IPv6 unicast
- VPNv4 unicast
- VPNv6 unicast

Task ID

Task ID	Operations
bgp	read, write

The following example shows how to schedule the delay of the BGP update generation in the start-up phase to prevent routes from being advertised to peers until RIB is synchronized.

```
Router# configure
Router(config)# router bgp 1
Router(config-bgp)# address-family ipv4 unicast
Router(config-bgp-af)# update wait-install delay startup 10
Router(config-bgp-af)# commit
```

use

To inherit configuration from a neighbor group, session group, or address family group, use the **use** command in an appropriate configuration mode. To discontinue inheritance from a group, use the **no** form of this command.

use {**af-group** *group-name* | **neighbor-group** *group-name* | **session-group** *group-nam* **e**}
no use {**af-group** [*group-name*] | **neighbor-group** [*group-name*] | **session-group** [*group-name*]}

Syntax Description

af-group	Specifies an address family group.
group-name	Name of the neighbor group, session group, or address family group from which you want to inherit configuration.
neighbor-group	Specifies a neighbor group.
session-group	Specifies a session group.

Command Default

Inheritance of group characteristics does not occur.

Command Modes

For **use af-group** version:

Address family group configuration

Neighbor address family configuration

Neighbor group address family configuration

For **use neighbor-group** version:

VRF neighbor configuration

For **use session-group** version:

Neighbor group configuration

Neighbor configuration

Session-group configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

The **use** command configures inheritance of configuration from an address family group, neighbor group, or session group, which means that any configuration for the group also takes effect for the user of the group.

The configuration inherited depends on the type of group that is specified. The group types are described in the following sections:

Address Family Group

An address family group can specify a configuration for only a single address family. The address family specified when the address family group was defined (through the **af-group** command) must match the address family from which the group is used.

Neighbor Group

A neighbor group (like a neighbor) can have address family-independent configuration and address family-specific configuration. All of these configurations could be inherited.

Session Group

A session group can have only address family-independent configuration and thus only address family-independent configuration is inherited from it.

The following rules govern inheritance to resolve possible conflicting configuration:

1. If a command is configured directly on the neighbor that is using group configuration, the command overrides the value that would be normally inherited from the group.
2. If the neighbor is configured to use a session group (for address family-independent configuration) or an address family group (for address family-specific configuration) and the command is configured for the session group or address family group, that configuration is used.
3. The neighbor group configuration is used:
   - If the command is not configured directly on the neighbor and the neighbor is not using a session group (for address family-independent configuration) or an af-group (for address family-specific configuration).
   - The neighbor is using a neighbor group and the command is configured on the neighbor group.

Typically, all configuration for a neighbor group is inherited, but some characteristics may be masked by a session group or address family group. For an example of this configuration, see the “Examples” section.

If the neighbor is using both a session group and a neighbor group and a specific command is configured for the neighbor group but not for the session group, then the configuration for the neighbor group does not take effect. The session group “hides” all address family-independent configuration on the neighbor group and prevents it from being inherited. Similarly, the use of an address family group hides any address family-specific configuration that may otherwise be inherited from a neighbor group for that address family.

In addition to neighbors using groups, it is possible to build a hierarchy by having groups use other groups. The following hierarchical groups are permitted:

- Session groups may use other session groups.
- Address family groups may use other address family groups.
- Neighbor groups may use session groups and address family groups.



Note Within the Cisco IOS XR system configuration architecture, do not combine the **remote-as** command and the **no use neighbor-group** command in the same commit, or the **remote-as** command and the **no use session-group** command in the same commit.

Task ID

Task ID	Operations
bgp	read, write

Examples

The following example shows how to define a session group session1 and configure neighbor 172.168.40.24 to use session1. As a result, the session1 configuration takes effect on the neighbor also.

```
RP/0/RP0/CPU0:router(config)# router bgp 1
RP/0/RP0/CPU0:router(config-bgp)# session-group session1
RP/0/RP0/CPU0:router(config-bgp-sngrp)# advertisement-interval 40
RP/0/RP0/CPU0:router(config-bgp-sngrp)# timers 30 90
RP/0/RP0/CPU0:router(config-bgp-sngrp)# exit
RP/0/RP0/CPU0:router(config-bgp)# neighbor 172.168.40.24
RP/0/RP0/CPU0:router(config-bgp-nbr)# remote-as 2
RP/0/RP0/CPU0:router(config-bgp-nbr)# use session-group session1
RP/0/RP0/CPU0:router(config-bgp-nbr)# exit
```

The following example is similar to the previous example, but in this case the **timers** command on the session group does not take effect on the neighbor because it is overridden by a **timers** command directly configured for the neighbor.

```
RP/0/RP0/CPU0:router(config)# router bgp 1
RP/0/RP0/CPU0:router(config-bgp)# session-group session1
RP/0/RP0/CPU0:router(config-bgp-sngrp)# advertisement-interval 40
RP/0/RP0/CPU0:router(config-bgp-sngrp)# timers 30 90
RP/0/RP0/CPU0:router(config-bgp-sngrp)# exit
RP/0/RP0/CPU0:router(config-bgp)# neighbor 172.168.40.24
RP/0/RP0/CPU0:router(config-bgp-nbr)# remote-as 2
RP/0/RP0/CPU0:router(config-bgp-nbr)# use session-group session1
RP/0/RP0/CPU0:router(config-bgp-nbr)# timers 60 180
RP/0/RP0/CPU0:router(config-bgp-nbr)# exit
```

The following example shows an address family group, family1, for IPv4 multicast and a neighbor group, neighbor1, that have IPv4 unicast and IPv4 multicast enabled. In this case, the neighbor inherits IPv4 unicast (and address family-independent) configuration from the neighbor group, but inherits IPv4 multicast configuration from the address family group. In this example, the neighbor group also has a remote autonomous system configured, so there is no need to configure a remote autonomous system for the neighbor because it inherits the remote autonomous system from the neighbor group:

```
RP/0/RP0/CPU0:router(config)# router bgp 1
RP/0/RP0/CPU0:router(config-bgp)# af-group family1 address-family ipv4 multicast
RP/0/RP0/CPU0:router(config-bgp-afgrp)# route-policy mcast-in in
RP/0/RP0/CPU0:router(config-bgp-afgrp)# exit
RP/0/RP0/CPU0:router(config-bgp)# neighbor-group neighbor1
RP/0/RP0/CPU0:router(config-bgp-nbrgrp)# remote-as 2
RP/0/RP0/CPU0:router(config-bgp-nbrgrp)# address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-bgp-nbrgrp-af)# route-policy policy1 in
RP/0/RP0/CPU0:router(config-bgp-nbrgrp-af)# route-policy policy1 out
RP/0/RP0/CPU0:router(config-bgp-nbrgrp-af)# exit
RP/0/RP0/CPU0:router(config-bgp-nbrgrp)# address-family ipv4 multicast
```

```
RP/0/RP0/CPU0:router(config-bgp-nbrgrp-af)# route-policy policy1 in
RP/0/RP0/CPU0:router(config-bgp-nbrgrp-af)# route-policy policy1 out
RP/0/RP0/CPU0:router(config-bgp-nbrgrp-af)# exit
RP/0/RP0/CPU0:router(config-bgp)# neighbor 172.168.40.24
RP/0/RP0/CPU0:router(config-bgp-nbr)# use neighbor-group neighbor1
RP/0/RP0/CPU0:router(config-bgp-nbr)# address-family ipv4 multicast
RP/0/RP0/CPU0:router(config-bgp-nbr-af)# use af-group family1
RP/0/RP0/CPU0:router(config-bgp-nbr-af)# exit
```

In the previous example, the neighbor uses the policy1 route policy for inbound and outbound IPv4 unicast routes, but uses the mcast-in route policy for inbound IPv4 multicast routes and no policy for outbound IPv4 multicast routes.

The following example shows a neighbor inheriting configuration from a session group that likewise inherits configuration from another session group. The configuration from both session groups take effect on the neighbor:

```
RP/0/RP0/CPU0:router(config)# router bgp 1
RP/0/RP0/CPU0:router(config-bgp)# session-group session1
RP/0/RP0/CPU0:router(config-bgp-sngrp)# advertisement-interval 40
RP/0/RP0/CPU0:router(config-bgp-sngrp)# exit
RP/0/RP0/CPU0:router(config-bgp)# session-group session2
RP/0/RP0/CPU0:router(config-bgp-sngrp)# use session-group session1
RP/0/RP0/CPU0:router(config-bgp-sngrp)# update-source Loopback0
RP/0/RP0/CPU0:router(config-bgp-sngrp)# exit
RP/0/RP0/CPU0:router(config-bgp)# neighbor 172.168.40.24
RP/0/RP0/CPU0:router(config-bgp-nbr)# remote-as 1
RP/0/RP0/CPU0:router(config-bgp-nbr)# use session-group session2
RP/0/RP0/CPU0:router(config-bgp-nbr)# exit
```

username (rpki-server)

To specify a SSH **username** for the RPKI cache-server, use the **username** command in rpki-server configuration mode. To remove the username, use the **no** form of this command.

username *user-name*
no username *user-name*

Syntax Description

user-name	Enters a username to be used for the SSH transport mechanism.

Command Default

Username is not configured.

Command Modes

RPKI server configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

The username configuration applies only if the SSH transport mechanism is active.

Task ID

Task ID	Operation
bgp	read, write

This example shows how to configure a username (*rpki-user*) for the RPKI cache-server SSH transport mechanism:

```
RP/0/RP0/CPU0:router#configure
RP/0/RP0/CPU0:router(config)#router bgp 100
RP/0/RP0/CPU0:router(config-bgp)#rpki server 172.168.35.40
RP/0/RP0/CPU0:router(config-bgp-rpki-cache)# transport ssh port 22
RP/0/RP0/CPU0:router(config-bgp-rpki-cache)#username rpki-user
```

vrf (BGP)

To configure a VPN routing and forwarding (VRF) instance and enter VRF configuration mode, use the **vrf** command in XR Config mode. To remove the VRF instance from the configuration file and restore the system to its default condition, use the **no** form of this command.



Note This feature is not supported.

vrf *vrf-name*
no vrf *vrf-name*

Syntax Description

vrf-name	Name of the VRF instance. The following names cannot be used: all, default, and global.

Command Default

No default behavior or values

Command Modes

XR Config mode

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

Use the **vrf** command to configure a VRF instance. A VRF instance is a collection of VPN routing and forwarding tables maintained at the provider edge (PE) router.

Task ID

Task ID	Operations
bgp	read, write

Examples

The following example shows how to configure a VRF instance and enter VRF configuration mode:

```
RP/0/RP0/CPU0:router(config)# router bgp 1
RP/0/RP0/CPU0:router(config-bgp)# vrf vrf-1
RP/0/RP0/CPU0:router(config-bgp-vrf)#
```

weight

To assign a weight to routes received from a neighbor, use the **weight** command in an appropriate configuration mode. To remove the **weight** command from the configuration file and restore the system to its default condition in which the software assigns the default weight to routes, use the **no** form of this command.

weight *weight-value*
no weight [*weight-value*]

Syntax Description

weight-value	Weight to assign. Range is 0 to 65535.

Command Default

Routes learned through another Border Gateway Protocol (BGP) peer have a default weight of 0 and routes sourced by the local router have a default weight of 32768.

Command Modes

IPv4 address family group configuration

IPv6 address family group configuration

IPv4 neighbor address family configuration

IPv4 neighbor group address family configuration

IPv6 neighbor group address family configuration

VPNv4 address family group configuration

VPNv4 neighbor address family configuration

VRF IPv4 neighbor address family configuration

VPNv4 neighbor group address family configuration

VPNv6 address family group configuration

VPNv6 neighbor address family configuration

VRF IPv6 neighbor address family configuration

VPNv6 neighbor group address family configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

The weight of a route is a Cisco-specific attribute. It is used in the best-path selection process (as the strongest tie-breaker). See the *Implementing BGP on* module of the *BGP Configuration Guide for Cisco NCS 5500 Series Routers* for information on best path. If there are two BGP routes with the same network layer reachability information (NLRI), the route with the higher weight is always chosen no matter what the value of other BGP attributes. Weight only has significance on the local router. Weight is assigned locally to the router, is a value that only makes sense to the specific router, is not propagated or carried through any route updates, and never is sent between BGP peers (even within the same AS).

Note If an address family group, neighbor group, or session group is configured, the configuration inside these configuration groups will not be effective unless it is applied directly or indirectly to one or more neighbors.

The weight assigned to individual routes can be further manipulated in the inbound route policy of a neighbor using the **set weight** command. The **set weight** command sets the weight directly. If you have particular neighbors that you want to prefer for most of your outbound traffic, you can assign a higher weight to all routes learned from that neighbor.

The weight assigned to individual routes may be modified by using an inbound routing policy.

Note For weight changes to take effect, you may need to use the clear bgp soft, on page 134 command.

If this command configures a neighbor group or neighbor address family group, all neighbors using the group inherit the configuration. Values of commands configured specifically for a neighbor override inherited values.

Task ID

Task ID	Operations
bgp	read, write

Examples

The following example shows how to assign a weight of 50 to all IP Version 4 (IPv4) unicast routes learned through 172.20.16.6:

```
RP/0/RP0/CPU0:router(config)# router bgp 1
RP/0/RP0/CPU0:router(config-bgp)# neighbor 172.20.16.6
RP/0/RP0/CPU0:router(config-bgp-nbr)# remote-as 1
RP/0/RP0/CPU0:router(config-bgp-nbr)# address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-bgp-nbr-af)# weight 50
RP/0/RP0/CPU0:router(config-bgp-nbr-af)# exit
```