

Routing Command Reference for Cisco ASR 9000 Series Routers

First Published: 2022-11-30

Last Modified: 2026-09-08

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

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

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

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

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

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

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

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

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

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

© 2022 Cisco Systems, Inc. All rights reserved.



CONTENTS

PREFACE

Preface xxxiii

Preface xxxiii

Communications, Services, and Additional Information xxxiii

CHAPTER 1

BGP Commands 1

accept-own 8

additional-paths install backup 9

additional-paths receive 10

additional-paths selection 12

additional-paths send 14

address-family (BGP) 16

advertise 19

advertise best-external 21

advertise local-labeled-route safi-unicast disable 23

advertise permanent-network 24

advertise vpn-imported label-mode per-vrf 25

advertisement-interval 26

af-group 28

aggregate-address 30

aigp 33

aigp send-cost-community 35

allocate-label 37

allow vpn default-originate 39

allowas-in 40

allowconfedas-in 41

as-format 42

as-league peers 43
as-override 45
as-path-loopcheck out disable 47
attribute-filter group 49
attribute large-community 50
bfd (BGP) 51
bgp as-path-loopcheck 57
bgp attribute-download 58
bgp auto-policy-soft-reset disable 60
bgp bestpath aigp ignore 61
bgp bestpath as-path ignore 62
bgp bestpath compare-routerid 64
bgp bestpath cost-community ignore 66
bgp bestpath med always 67
bgp bestpath med confed 69
bgp bestpath med missing-as-worst 71
bgp bestpath origin-as allow invalid 73
bgp bestpath origin-as use validity 75
bgp bestpath as-path multipath-relax 76
bestpath as-path multipath-relax (BGP IPv4 RT-filter address family) 77
bgp client-to-client reflection disable 78
bgp cluster-id 80
bgp confederation identifier 82
bgp confederation peers 84
bgp dampening 86
bgp default local-preference 88
bgp enforce-first-as disable 89
bgp fast-external-fallover disable 90
bgp graceful-restart 91
bgp graceful-restart graceful-reset 93
bgp graceful-restart purge-time 94
bgp graceful-restart restart-time 95
bgp graceful-restart stalepath-time 96
bgp import-delay 98

bgp label-delay 99
bgp log neighbor changes disable 101
bgp lpts-secure-binding 103
bgp maximum neighbor 104
bgp multipath as-path 105
bgp origin-as validation enable 106
bgp policy propagation input flow-tag 107
bgp redistribute-internal 108
bgp router-id 110
bgp scan-time 111
bgp unsafe-ebgp-policy 113
bgp update-delay 114
bgp write-limit 116
bind-source (rpki-server) 118
bmp-activate 119
bmp advertisement-interval 120
bmp scan-time 122
bmp server 123
capability additional-paths receive 126
capability additional-paths send 128
capability orf prefix 130
capability suppress 132
cef consistency-hashing auto-recovery 135
cef fast-reroute follow bgp-pic 136
clear bgp 137
clear bgp current-mode 139
clear bgp dampening 141
clear bgp external 143
clear bgp flap-statistics 145
clear bgp long-lived-stale 147
clear bgp nexthop performance-statistics 148
clear bgp nexthop registration 150
clear bgp peer-drops 152
clear bgp performance-statistics 153

clear bgp rpki 154
clear bgp self-originated 156
clear bgp shutdown 158
clear bgp soft 160
cluster-id allow-equal 163
default-information originate (BGP) 164
default-martian-check disable 165
default-metric (BGP) 166
default-originate 167
description (BGP) 169
distance bgp 170
distribute bgp-ls (ISIS) 172
distribute bgp-ls (OSPF) 173
distributed speaker 174
dmz-link-bandwidth 175
domain-distinguisher 177
dscp (BGP) 178
dynamic-med interval 180
ebgp-multihop 181
ecmp-delay 183
export route-policy 184
export route-target 185
graceful-maintenance 187
host-reachability protocol bgp 189
ibgp policy out enforce-modifications 190
ignore-connected-check 192
import 193
import route-policy 195
import route-target 196
keychain 198
keychain-disable 200
keychain inheritance-disable 202
label-allocation-mode 204
label mode 206

local-as 210
long-lived-graceful-restart 212
lpts punt excessive-flow-trap penalty-timeout bgp 214
lpts punt excessive-flow-trap routing-protocols-enable 215
match flow-tag 216
maximum-paths (BGP) 218
maximum-paths external-rtfilter (BGP IPv4 RT-filter address family) 220
maximum-prefix (BGP) 221
mpls activate (BGP) 224
multipath 227
mvpn 228
neighbor (BGP) 229
neighbor-group 231
neighbor internal-vpn-client 233
network backdoor 234
network (BGP) 236
next-hop-self 238
next-hop-unchanged 240
bgp nexthop resolution allow-default 242
nexthop resolution prefix-length minimum 243
nexthop route-policy 245
nexthop trigger-delay 247
nsr (BGP) 249
nsr disable (BGP) 251
optimal-route-reflection 253
orf 255
password (BGP) 257
password-disable 259
password (rpki-server) 261
permanent-network 262
precedence 263
preference (rpki-server) 265
purge-time (rpki-server) 266
rd 267

receive-buffer-size 269
redistribute (BGP) 271
refresh-time (rpki-server) 276
remote-as (BGP) 277
remove-private-as 280
replace as-path all 282
response-time (rpki-server) 284
retain local-label 285
retain route-target 286
route-monitoring inbound post-policy 288
route-monitoring local-rib 289
route-policy (BGP) 290
route-reflector-client 293
router bgp 295
rpki route 296
rpki server 298
selective-vrf-download disable 299
send-buffer-size 300
send-community-ebgp 302
send-community-gshut-ebgp 304
send-extended-community-ebgp 305
session-group 307
session-open-mode 309
set flow-tag 311
show bgp 312
show bgp advertised 327
show bgp as-neighbors 334
show bgp af-group 335
show bgp aggregate-contributors 338
show bgp attribute-key 340
show bgp bmp 344
show bgp cidr-only 370
show bgp community 374
show bgp convergence 379

show bgp dampened-paths 382
show bgp flap-statistics 386
show bgp inconsistent-as 391
show bgp l2vpn 395
show bgp l2vpn vpls 399
show bgp labels 405
show bgp large-community 408
show bgp memory history 411
show bgp neighbor-group 413
show bgp neighbors 417
show bgp neighbors nsr 436
show bgp nexthops 438
show bgp nsr 447
show bgp paths 451
show bgp policy 454
show bgp process 462
show bgp regexp 482
show bgp route-policy 486
show bgp session-group 490
show bgp sessions 493
show bgp summary 496
show bgp summary nsr 500
show bgp table 504
show bgp truncated-communities 507
show bgp update-group 511
show bgp update in error process 518
show bgp update out 519
show bgp update out filter-group 523
show bgp update out process 524
show bgp update out sub-group 526
show bgp update out update-group 528
show bgp vrf 530
show bgp vrf update in error 533
show lpts punt excessive-flow-trap bgp 534

show protocols (BGP) 535
show running-config lpts punt excessive-flow-trap 537
show svd role 539
show svd state 540
show tcp pcb 541
shutdown (BGP) 543
shutdown (rpki-server) 545
signalling disable 546
site-of-origin (BGP) 547
slow peer (BGP neighbor address-family configuration) 549
slow peer (BGP router global configuration) 552
socket receive-buffer-size 554
socket send-buffer-size 556
soft-reconfiguration inbound 558
speaker-id 561
svd platform enable 562
table-policy 564
tcp mss inheritance-disable 566
timers (BGP) 567
timers bgp 569
transport (rpki-server) 571
ttl-security 573
update in error-handling basic disable 576
update in error-handling extended 577
update in error-handling treat-as-withdraw (BGP) 578
update in labeled-unicast equivalent 579
update limit 580
update limit address-family 581
update limit sub-group 583
update out logging 585
update out originator-loopcheck disable 586
update-source 588
update wait-install 590
update wait-install delay startup 591

use 592
username (rpki-server) 597
vrf (BGP) 598
weight 599

CHAPTER 2 BGP Flowspec Commands 601
class-map type traffic (BGP-flowspec) 602
class type traffic 603
destination-prefix in 604
drop (BGP-flowspec) 606
flowspec 607
flowspec disable 608
local-install 609
match destination-address 610
match destination-port 611
match fragment-type 612
match icmp code 613
match icmp type 614
match packet length 615
match source-address 616
match source-port 618
match tcp flag 619
redirect (BGP Flowspec) 620
service-policy 621
show flowspec 622
source prefix 623

CHAPTER 3 BFD Commands 625
address-family ipv4 unicast (BFD) 627
bfd 629
bfd address-family ipv4 destination 631
bfd address-family ipv4 echo minimum-interval 632
bfd address-family ipv4 fast-detect 634
bfd address-family ipv4 minimum-interval 635

bfd address-family ipv4 multiplier 638
bfd address-family ipv4 timers 641
bundle coexistence bob-blb 643
bfd dampening 645
bfd dampening disable 648
bfd echo ipv4 bundle-per-member minimum-interval 649
bfd encap-mode 651
bfd fast-detect 652
bfd minimum-interval 655
bfd mode 659
bfd multipath include location 660
bfd multiplier 662
bfd bundle per-member mode 665
clear bfd counters 666
clear bfd dampening 668
echo disable 670
echo ipv4 source 672
echo latency detect 674
echo startup validate 676
hw-module bfd-hw-offload 678
interface (BFD) 680
ipv6 checksum 683
multihop ttl-drop-threshold 685
show bfd 687
show bfd client 690
show bfd counters 692
show bfd hw-offload 695
show bfd mib session 697
show bfd multipath 700
show bfd session 702
show bfd summary 708

CHAPTER 4 **EIGRP Commands** 711
address-family (EIGRP) 713

authentication keychain 715
auto-summary (EIGRP) 717
autonomous-system 719
bandwidth-percent (EIGRP) 720
bfd fast-detect (EIGRP) 721
bfd minimum-interval (EIGRP) 722
bfd multiplier (EIGRP) 724
clear eigrp neighbors 725
clear eigrp topology 727
default-information 729
default-metric (EIGRP) 730
distance (EIGRP) 732
hello-interval (EIGRP) 734
hold-time (EIGRP) 735
interface (EIGRP) 737
log-neighbor-changes 739
log-neighbor-warnings 740
maximum-paths (EIGRP) 741
maximum-prefix (EIGRP) 742
metric (EIGRP) 744
metric maximum-hops 746
metric rib-scale 747
metric version (EIGRP) 748
metric weights 749
neighbor (EIGRP) 751
neighbor maximum-prefix 753
next-hop-self disable 756
nsf disable (EIGRP) 757
passive-interface (EIGRP) 759
redistribute (EIGRP) 760
redistribute maximum-prefix 762
remote-neighbor (unicast) 764
route-policy (EIGRP) 766
router eigrp 767

router-id (EIGRP) 768
show eigrp accounting 769
show eigrp interfaces 771
show eigrp neighbors 775
show eigrp topology 778
show eigrp traffic 781
show protocols (EIGRP) 783
site-of-origin (EIGRP) 786
split-horizon disable (EIGRP) 788
stub (EIGRP) 789
summary-address (EIGRP) 791
timers active-time 793
timers nsf route-hold (EIGRP) 794
variance 795
vrf (EIGRP) 796

CHAPTER 5 **IS-IS Commands** 799
address-family (IS-IS) 803
address-family multicast topology (IS-IS) 805
adjacency-check disable 807
advertise passive-only 809
adjacency stagger (IS-IS) 810
algorithm 0 812
attached-bit receive ignore 813
attached-bit send 814
authentication-check disable 816
circuit-type 817
clear isis process 819
clear isis route 820
clear isis statistics 822
cost-fallback anomaly delay 824
csnp-interval 826
default-information originate (IS-IS) 828
disable (IS-IS) 830

distance (IS-IS) **831**
distribute-list in (IS-IS) **833**
fast-reroute per-link (IS-IS) **835**
fast-reroute per-prefix (IS-IS) **837**
fast-reroute per-link priority-limit (IS-IS) **839**
fast-reroute per-link use-candidate-only (IS-IS) **841**
fast-reroute per-prefix load-sharing disable (IS-IS) **842**
fast-reroute per-prefix tiebreaker (IS-IS) **843**
fast-reroute per-prefix use-candidate-only (IS-IS) **845**
hello-interval (IS-IS) **846**
hello-multiplier **848**
hello-padding **850**
hello-password **852**
hello-password keychain **854**
hello-password accept **856**
hostname dynamic disable **858**
ignore-lsp-errors **859**
instance-id **860**
interface (IS-IS) **861**
ipfrr lfa **863**
ipfrr lfa exclude interface **865**
ispf **867**
is-type **868**
link-down fast-detect **870**
link-group **871**
log adjacency changes (IS-IS) **872**
log pdu drops **873**
lsp fast-flood threshold **874**
lsp-fast-flooding **875**
lsp-gen-interval **877**
lsp-interval **879**
lsp-mtu **880**
lsp-password **882**
lsp-password accept **884**

lsp-refresh-interval 886
maximum-paths (IS-IS) 888
maximum-redistributed-prefixes (IS-IS) 889
maximum-propagated-prefixes (IS-IS) 890
max-lsp-lifetime 891
max-lsp-tx 892
max-metric 893
metric fallback anomaly 895
mesh-group (IS-IS) 897
metric (IS-IS) 899
metric-style narrow 902
metric-style transition 904
metric-style wide 906
microloop avoidance (IS-IS) 908
min-lsp-arrivaltime 910
mpls ldp auto-config 912
mpls ldp sync (IS-IS) 913
mpls traffic-eng (IS-IS) 915
mpls traffic-eng multicast-intact (IS-IS) 917
mpls traffic-eng path-selection ignore overload 918
mpls traffic-eng router-id (IS-IS) 920
net 922
nsf (IS-IS) 924
nsf interface-expires 926
nsf interface-timer 928
nsf lifetime (IS-IS) 930
oor-set-overload-bit disable 931
passive (IS-IS) 932
point-to-point 933
priority (IS-IS) 934
propagate level 936
protocol shutdown (IS-IS) 938
psnp-interval 939
redistribute (IS-IS) 940

remote-psnp-delay 944
retransmit-interval (IS-IS) 946
retransmit-throttle-interval 948
router isis 950
router isis static neighbor 951
route source first-hop 952
set default-isis-instance 953
set-overload-bit 955
set-attached-bit 957
show isis 959
show isis adjacency 961
show isis adjacency-log 963
show isis checkpoint adjacency 965
show isis checkpoint interface 967
show isis checkpoint lsp 969
show isis database 971
show isis database-log 972
show isis fast-reroute 974
show isis hostname 977
show isis interface 979
show isis lsp-log 983
show isis mesh-group 985
show isis mpls traffic-eng adjacency-log 987
show isis mpls traffic-eng advertisements 989
show isis neighbors 992
show isis protocol 995
show isis route 998
show isis spf-log 1001
show isis statistics 1008
show isis topology 1012
show protocols (IS-IS) 1015
shutdown (IS-IS) 1019
single-topology 1020
snmp-server traps isis 1021

spf-interval 1022
spf-interval ietf 1024
spf prefix-priority (IS-IS) 1026
summary-prefix (IS-IS) 1028
suppressed 1030
tag (IS-IS) 1031
topology-id 1032
trace (IS-IS) 1033
ucmp (IS-IS) 1034
ucmp delay-interval 1036
ucmp exclude interface 1038

CHAPTER 6 OSPF Commands 1041
address-family (OSPF) 1044
adjacency stagger (OSPF) 1045
area (OSPF) 1047
authentication (OSPF) 1049
authentication-key (OSPF) 1051
auto-cost (OSPF) 1053
capability opaque disable 1055
capability vrf-lite (OSPFv2) 1056
clear ospf process 1057
clear ospf redistribution 1059
clear ospf routes 1061
clear ospf statistics 1062
cost (OSPF) 1064
cost-fallback (OSPF) 1066
database-filter all out (OSPF) 1068
dead-interval (OSPF) 1069
default-cost (OSPF) 1071
default-information originate (OSPF) 1073
default-metric (OSPF) 1075
demand-circuit (OSPF) 1077
disable-dn-bit-check 1079

distance (OSPF) **1080**
distance ospf **1083**
distribute-list **1085**
domain-id (OSPF) **1087**
domain-tag **1089**
exchange-timer **1090**
fast-reroute (OSPFv2) **1092**
fast-reroute per-link exclude interface **1094**
fast-reroute per-prefix exclude interface (OSPFv2) **1096**
fast-reroute per-prefix lfa-candidate (OSPFv2) **1098**
fast-reroute per-prefix remote-lfa (OSPFv2) **1099**
fast-reroute per-prefix ti-lfa **1101**
fast-reroute per-prefix use-candidate-only (OSPFv2) **1102**
flood-reduction (OSPF) **1103**
hello-interval (OSPF) **1105**
ignore lsa mospf **1107**
interface (OSPF) **1108**
log adjacency changes (OSPF) **1110**
loopback stub-network **1111**
lpts punt excessive-flow-trap penalty-timeout ospf **1112**
max external-lsa **1113**
max-lsa **1115**
max-metric **1118**
maximum interfaces (OSPF) **1121**
maximum paths (OSPF) **1123**
maximum redistributed-prefixes (OSPF) **1125**
message-digest-key **1127**
microloop avoidance (OSPF) **1130**
mpls ldp auto-config (OSPF) **1131**
mpls ldp sync (OSPF) **1132**
mpls traffic-eng (OSPF) **1133**
mpls traffic-eng igp-intact (OSPF) **1135**
mpls traffic-eng multicast-intact (OSPF) **1137**
mpls traffic-eng ldp-sync-update (OSPF) **1138**

mpls traffic-eng router-id (OSPF) 1139
mtu-ignore (OSPF) 1141
multi-area-interface 1143
neighbor (OSPF) 1145
neighbor database-filter all out 1147
network (OSPF) 1148
nsf (OSPF) 1150
nsf flush-delay-time (OSPF) 1152
nsf interval (OSPF) 1153
nsf lifetime (OSPF) 1154
nsr (OSPF) 1155
nssa (OSPF) 1157
ospf name-lookup 1159
packet-size (OSPF) 1160
passive (OSPF) 1162
prefix-suppression 1164
priority (OSPF) 1166
protocol shutdown 1168
queue dispatch flush-lsa 1169
queue dispatch incoming 1171
queue dispatch rate-limited-lsa 1173
queue dispatch spf-lsa-limit 1175
queue limit 1176
range (OSPF) 1178
redistribute (OSPF) 1180
retransmit-interval (OSPF) 1185
route-policy (OSPF) 1187
router-id (OSPF) 1188
router ospf (OSPF) 1190
security ttl (OSPF) 1192
segment-routing sr-prefer prefix-list 1194
sham-link 1196
show lpts punt excessive-flow-trap ospf 1198
show ospf 1199

show ospf border-routers 1202
show ospf database 1204
show ospf flood-list 1219
show ospf interface 1221
show ospf mpls traffic-eng 1224
show ospf message-queue 1229
show ospf neighbor 1232
show ospf request-list 1241
show ospf retransmission-list 1244
show ospf routes 1246
show ospf sham-links 1251
show ospf summary-prefix 1254
show ospf virtual-links 1256
show protocols (OSPF) 1258
snmp context (OSPF) 1261
snmp trap (OSPF) 1263
snmp trap rate-limit (OSPF) 1264
spf prefix-priority (OSPFv2) 1265
stub (OSPF) 1267
summary-prefix (OSPF) 1269
timers lsa group-pacing 1271
timers lsa min-arrival 1272
timers throttle lsa all (OSPF) 1273
timers throttle spf (OSPF) 1276
transmit-delay (OSPF) 1278
virtual-link (OSPF) 1280
vrf (OSPF) 1282

CHAPTER 7 **OSPFv3 Commands** 1285

address-family (OSPFv3) 1288
area (OSPFv3) 1289
authentication (OSPFv3) 1290
auto-cost (OSPFv3) 1292
capability vrf-lite (OSPFv3) 1294

clear ospfv3 process 1296
clear ospfv3 redistribution 1298
clear ospfv3 routes 1299
clear ospfv3 statistics 1300
cost (OSPFv3) 1302
database-filter all out (OSPFv3) 1304
dead-interval (OSPFv3) 1306
default-cost (OSPFv3) 1308
default-information originate (OSPFv3) 1310
default-metric (OSPFv3) 1312
demand-circuit (OSPFv3) 1314
distance ospfv3 1316
distribute-list prefix-list in 1318
distribute-list prefix-list out 1320
domain-id (OSPFv3) 1322
encryption 1324
flood-reduction (OSPFv3) 1326
graceful-restart (OSPFv3) 1328
hello-interval (OSPFv3) 1329
instance (OSPFv3) 1331
interface (OSPFv3) 1333
log adjacency changes (OSPFv3) 1335
maximum interfaces (OSPFv3) 1336
maximum paths (OSPFv3) 1337
maximum redistributed-prefixes (OSPFv3) 1338
mtu-ignore (OSPFv3) 1340
neighbor (OSPFv3) 1342
network (OSPFv3) 1344
nssa (OSPFv3) 1346
nsr (OSPFv3) 1348
ospfv3 name-lookup 1350
packet-size (OSPFv3) 1351
passive (OSPFv3) 1352
priority (OSPFv3) 1354

protocol shutdown (OSPFv3) 1356
range (OSPFv3) 1357
redistribute (OSPFv3) 1359
retransmit-interval (OSPFv3) 1363
router-id (OSPFv3) 1365
router ospfv3 1367
sham-link (OSPFv3) 1368
show ospfv3 1370
show ospfv3 border-routers 1373
show ospfv3 database 1375
show ospfv3 flood-list 1387
show ospfv3 interface 1389
show ospfv3 message-queue 1392
show ospfv3 neighbor 1394
show ospfv3 request-list 1400
show ospfv3 retransmission-list 1402
show ospfv3 routes 1404
show ospfv3 statistics rib-thread 1406
show ospfv3 summary-prefix 1408
show ospfv3 virtual-links 1410
show protocols (OSPFv3) 1412
snmp context (OSPFv3) 1414
snmp trap (OSPFv3) 1416
snmp trap rate-limit (OSPFv3) 1417
spf prefix-priority (OSPFv3) 1418
stub (OSPFv3) 1420
stub-router 1422
summary-prefix (OSPFv3) 1424
timers lsa arrival 1426
timers pacing flood 1427
timers pacing lsa-group 1429
timers pacing retransmission 1431
timers throttle lsa all (OSPFv3) 1433
timers throttle spf (OSPFv3) 1435

trace (OSPFv3) 1437
transmit-delay (OSPFv3) 1439
virtual-link (OSPFv3) 1441
vrf (OSPFv3) 1443

CHAPTER 8 RIB Commands 1445
address-family next-hop dampening disable 1447
clear route 1448
maximum prefix (RIB) 1450
lcc 1452
rcc 1453
recursion-depth-max 1455
router rib 1456
rump always-replicate 1457
show lcc statistics 1458
show mpls lsd forwarding 1460
show rcc 1461
show rcc statistics 1463
show rcc vrf 1465
show rib 1466
show rib afi-all 1468
show rib attributes 1470
show rib client-id 1471
show rib clients 1473
show rib extcomms 1475
show rib firsthop 1477
show rib history 1479
show rib next-hop 1481
show rib opaques 1483
show rib protocols 1485
show rib recursion-depth-max 1487
show rib statistics 1489
show rib tables 1492
show rib trace 1494

show rib vpn-attributes 1496
show rib vrf 1498
show route 1500
show route backup 1506
show route best-local 1509
show route connected 1511
show route local 1513
show route longer-prefixes 1515
show route next-hop 1517
show route quarantined 1519
show route resolving-next-hop 1521
show route static 1523
show route summary 1525

CHAPTER 9 RIP Commands 1529
authentication keychain mode (RIP) 1531
auto-summary (RIP) 1533
broadcast-for-v2 1534
clear rip 1535
clear rip database 1537
clear rip interface 1539
clear rip out-of-memory 1541
clear rip statistics 1543
default-information originate (RIP) 1545
default-metric (RIP) 1546
distance (RIP) 1548
interface (RIP) 1551
maximum-paths (RIP) 1553
metric-zero-accept 1554
neighbor (RIP) 1555
nsf (RIP) 1557
output-delay 1558
passive-interface (RIP) 1559
poison-reverse 1561

receive version 1563
redistribute (RIP) 1564
router rip 1568
route-policy (RIP) 1570
send version 1572
show protocols (RIP) 1573
show rip 1575
show rip database 1577
show rip interface 1580
show rip statistics 1583
site-of-origin (RIP) 1585
split-horizon disable (RIP) 1587
timers basic 1589
validate-update-source disable 1591
vrf (RIP) 1592

CHAPTER 10 **Routing Policy Language Commands** 1595
abort (RPL) 1599
add 1601
apply 1603
as-path in 1605
as-path is-local 1607
as-path length 1609
as-path neighbor-is 1611
as-path originates-from 1614
as-path passes-through 1616
as-set 1618
as-path-set 1620
as-path unique-length 1622
community is-empty 1624
community matches-any 1625
community matches-every 1627
community matches-within 1629
community-set 1631

delete community 1634
delete large-community 1636
delete extcommunity rt 1638
destination in 1640
done 1642
drop 1644
edit 1646
end-global 1649
end-policy 1650
end-set 1651
event prefix route-policy pass 1653
extcommunity rt is-empty 1654
extcommunity rt matches-any 1655
extcommunity rt matches-every 1657
extcommunity rt matches-within 1659
extcommunity-set bandwidth 1661
extcommunity-set cost 1662
extcommunity-set opaque 1664
extcommunity-set rt 1665
extcommunity-set seg-nh 1667
extcommunity-set soo 1668
extcommunity soo is-empty 1670
extcommunity soo matches-any 1671
extcommunity soo matches-every 1673
globalVarN is 1675
if 1677
if route-aggregated 1683
is-best-path 1684
is-backup-path 1685
is-multi-path 1686
local-preference 1687
large-community is-empty 1688
large-community matches-any 1690
large-community matches-every 1692

large-community matches-within 1694
large-community-set 1696
med 1697
next-hop in 1699
orf prefix in 1701
origin is 1703
ospf-area 1705
ospf-area-set 1707
pass 1709
path-type is 1711
policy-global 1712
prefix-set 1714
prepend as-path 1716
protocol 1718
rd in 1720
rd-set 1721
replace as-path 1723
remove as-path private-as 1725
rib-has-route 1726
route-has-label 1728
route-policy (RPL) 1729
route-type is 1731
rpl editor 1733
rpl maximum 1735
rpl set-exit-as-abort 1737
set administrative-distance 1738
set aigp-metric 1739
set community 1740
set core-tree 1742
set dampening 1743
set fallback-vrf-lookup 1745
set eigrp-metric 1746
set extcommunity color 1748
set extcommunity cost 1749

set extcommunity rt 1751
set ip-precedence 1753
set isis-metric 1755
set label 1756
set label-mode 1757
set large-community 1759
set level 1761
set local-preference 1762
set med 1763
set metric-type (IS-IS) 1765
set metric-type (OSPF) 1766
set next-hop 1767
set origin 1769
set ospf-metric 1770
set path-selection 1771
set qos-group (RPL) 1773
set rib-metric 1774
set rip-metric 1775
set rip-tag 1776
set rpf-topology 1777
set rtset route-limit 1779
set spf-priority 1780
set tag 1781
set traffic-index 1782
set vpn-distinguisher 1784
set weight 1785
show rpl 1787
show rpl active as-path-set 1789
show rpl active community-set 1792
show rpl active extcommunity-set 1795
show rpl active prefix-set 1798
show rpl active rd-set 1801
show rpl active route-policy 1803
show rpl as-path-set 1805

show rpl as-path-set attachpoints 1807
show rpl as-path-set references 1810
show rpl community-set 1812
show rpl community-set attachpoints 1814
show rpl community-set references 1816
show rpl extcommunity-set 1819
show rpl inactive as-path-set 1822
show rpl inactive community-set 1825
show rpl inactive extcommunity-set 1828
show rpl inactive prefix-set 1831
show rpl inactive rd-set 1834
show rpl inactive route-policy 1836
show rpl maximum 1839
show rpl policy-global references 1841
show rpl prefix-set 1843
show rpl prefix-set attachpoints 1845
show rpl prefix-set references 1848
show rpl rd-set 1850
show rpl rd-set attachpoints 1852
show rpl rd-set references 1854
show rpl route-policy 1856
show rpl route-policy attachpoints 1859
show rpl route-policy inline 1862
show rpl route-policy references 1864
show rpl route-policy uses 1867
show rpl unused as-path-set 1870
show rpl unused community-set 1873
show rpl unused extcommunity-set 1876
show rpl unused prefix-set 1878
show rpl unused rd-set 1881
show rpl unused route-policy 1883
source in 1886
source rt 1888
suppress-route 1889

tag 1890
tag in 1891
tag-set 1893
unsuppress-route 1894
var globalVarN 1896
vpn-distinguisher is 1897

CHAPTER 11 Static Routing Commands 1899
address-family (static) 1900
maximum path (static) 1902
metric (static) 1904
route (static) 1906
router static 1910
track (static) 1912
vrf (static) 1914

CHAPTER 12 RCMD Commands 1917
router-convergence 1918
monitor-convergence (IS-IS) 1919
monitor-convergence (OSPF) 1920
collect-diagnostics (RCMD) 1921
event-buffer-size (RCMD) 1923
max-events-stored (RCMD) 1924
monitoring-interval (RCMD) 1925
node disable (RCMD) 1927
prefix-list (monitor-convergence IS-IS) 1929
prefix-list (monitor-convergence OSPF) 1931
priority (RCMD) 1933
protocol (RCMD) 1935
show rcmd isis event prefix 1936
show rcmd ospf event prefix 1938
show rcmd ospf event spf 1940
storage-location 1943
track-external-routes 1945

track-summary-routes 1946

CHAPTER 13 **Locator/ID Separation Protocol Commands** 1947

Locator/ID Separation Protocol on Cisco IOS XR 1948

address-family (LISP) 1949

clear lisp vrf 1951

decapsulation filter rloc source 1952

eid-mtu 1954

eid-table 1956

etr 1958

etr accept-map-request-mapping 1959

etr map-cache-ttl 1961

etr map-server 1963

itr map-resolver 1965

locator reachability 1967

locator-set 1968

locator-table 1970

loc-reach-algorithm rloc-probing 1972

map-cache-limit 1974

map-cache 1975

map-request-source 1977

map-server rloc members distribute 1978

map-server rloc members modify-discovered {add | override} 1980

other-xtr-probe 1982

proxy-etr 1984

proxy-itr 1986

remote-rloc-probe 1988

router lisp 1989

show lisp decapsulation filter 1991

show lisp session 1992

show lisp site rloc members 1993

show lisp site 1994

solicit-map-request 1996

use-petr 1998



Preface

The *Routing Command Reference for Cisco ASR 9000 Series Routers* preface contains these sections:

- Preface, on page xxxiii
- Communications, Services, and Additional Information, on page xxxiii

Preface

Initial release of the cumulative command reference document that covers all updates from Release 3.9.0 onwards.

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

This chapter describes the commands used to configure and monitor Border Gateway Protocol (BGP) on Cisco ASR 9000 Series Aggregation Services 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 *Routing Configuration Guide for Cisco ASR 9000 Series Routers*.

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.

Note Use only supported CLI commands for native CLI configurations. Avoid commands marked `Deprecated` in the CLI help output for the indicated release and later releases. These deprecated commands remain available only to support OpenConfig configurations and are not intended for native CLI use.

- accept-own, on page 8
- additional-paths install backup, on page 9
- additional-paths receive, on page 10
- additional-paths selection, on page 12
- additional-paths send, on page 14
- address-family (BGP), on page 16
- advertise, on page 19
- advertise best-external, on page 21
- advertise local-labeled-route safi-unicast disable, on page 23
- advertise permanent-network, on page 24
- advertise vpn-imported label-mode per-vrf, on page 25
- advertisement-interval, on page 26
- af-group, on page 28
- aggregate-address, on page 30
- aigp, on page 33

- aigp send-cost-community, on page 35
- allocate-label, on page 37
- allow vpn default-originate, on page 39
- allowas-in, on page 40
- allowconfedas-in, on page 41
- as-format, on page 42
- as-league peers, on page 43
- as-override, on page 45
- as-path-loopcheck out disable, on page 47
- attribute-filter group, on page 49
- attribute large-community, on page 50
- bfd (BGP), on page 51
- bgp as-path-loopcheck, on page 57
- bgp attribute-download, on page 58
- bgp auto-policy-soft-reset disable, on page 60
- bgp bestpath aigp ignore, on page 61
- bgp bestpath as-path ignore, on page 62
- bgp bestpath compare-routerid, on page 64
- bgp bestpath cost-community ignore, on page 66
- bgp bestpath med always, on page 67
- bgp bestpath med confed, on page 69
- bgp bestpath med missing-as-worst, on page 71
- bgp bestpath origin-as allow invalid, on page 73
- bgp bestpath origin-as use validity, on page 75
- bgp bestpath as-path multipath-relax , on page 76
- bestpath as-path multipath-relax (BGP IPv4 RT-filter address family), on page 77
- bgp client-to-client reflection disable, on page 78
- bgp cluster-id, on page 80
- bgp confederation identifier, on page 82
- bgp confederation peers, on page 84
- bgp dampening, on page 86
- bgp default local-preference, on page 88
- bgp enforce-first-as disable, on page 89
- bgp fast-external-fallover disable, on page 90
- bgp graceful-restart, on page 91
- bgp graceful-restart graceful-reset, on page 93
- bgp graceful-restart purge-time, on page 94
- bgp graceful-restart restart-time, on page 95
- bgp graceful-restart stalepath-time, on page 96
- bgp import-delay, on page 98
- bgp label-delay, on page 99
- bgp log neighbor changes disable, on page 101
- bgp lpts-secure-binding, on page 103
- bgp maximum neighbor, on page 104
- bgp multipath as-path, on page 105
- bgp origin-as validation enable, on page 106

- bgp policy propagation input flow-tag, on page 107
- bgp redistribute-internal, on page 108
- bgp router-id, on page 110
- bgp scan-time, on page 111
- bgp unsafe-ebgp-policy, on page 113
- bgp update-delay, on page 114
- bgp write-limit, on page 116
- bind-source (rpki-server), on page 118
- bmp-activate, on page 119
- bmp advertisement-interval, on page 120
- bmp scan-time, on page 122
- bmp server, on page 123
- capability additional-paths receive, on page 126
- capability additional-paths send, on page 128
- capability orf prefix, on page 130
- capability suppress , on page 132
- cef consistency-hashing auto-recovery, on page 135
- cef fast-reroute follow bgp-pic, on page 136
- clear bgp, on page 137
- clear bgp current-mode, on page 139
- clear bgp dampening, on page 141
- clear bgp external, on page 143
- clear bgp flap-statistics, on page 145
- clear bgp long-lived-stale, on page 147
- clear bgp nexthop performance-statistics, on page 148
- clear bgp nexthop registration, on page 150
- clear bgp peer-drops, on page 152
- clear bgp performance-statistics, on page 153
- clear bgp rpki , on page 154
- clear bgp self-originated, on page 156
- clear bgp shutdown, on page 158
- clear bgp soft, on page 160
- cluster-id allow-equal , on page 163
- default-information originate (BGP), on page 164
- default-martian-check disable, on page 165
- default-metric (BGP), on page 166
- default-originate, on page 167
- description (BGP), on page 169
- distance bgp, on page 170
- distribute bgp-ls (ISIS), on page 172
- distribute bgp-ls (OSPF), on page 173
- distributed speaker, on page 174
- dmz-link-bandwidth, on page 175
- domain-distinguisher, on page 177
- dscp (BGP), on page 178
- dynamic-med interval, on page 180

- ebgp-multihop, on page 181
- ecmp-delay, on page 183
- export route-policy, on page 184
- export route-target, on page 185
- graceful-maintenance, on page 187
- host-reachability protocol bgp, on page 189
- ibgp policy out enforce-modifications, on page 190
- ignore-connected-check, on page 192
- import, on page 193
- import route-policy, on page 195
- import route-target, on page 196
- keychain, on page 198
- keychain-disable, on page 200
- keychain inheritance-disable, on page 202
- label-allocation-mode, on page 204
- label mode, on page 206
- local-as, on page 210
- long-lived-graceful-restart, on page 212
- lpts punt excessive-flow-trap penalty-timeout bgp, on page 214
- lpts punt excessive-flow-trap routing-protocols-enable, on page 215
- match flow-tag, on page 216
- maximum-paths (BGP), on page 218
- maximum-paths external-rtfilter (BGP IPv4 RT-filter address family), on page 220
- maximum-prefix (BGP), on page 221
- mpls activate (BGP), on page 224
- multipath, on page 227
- mvpn, on page 228
- neighbor (BGP), on page 229
- neighbor-group, on page 231
- neighbor internal-vpn-client , on page 233
- network backdoor, on page 234
- network (BGP), on page 236
- next-hop-self, on page 238
- next-hop-unchanged, on page 240
- bgp nexthop resolution allow-default, on page 242
- nexthop resolution prefix-length minimum, on page 243
- nexthop route-policy, on page 245
- nexthop trigger-delay, on page 247
- nsr (BGP), on page 249
- nsr disable (BGP), on page 251
- optimal-route-reflection, on page 253
- orf, on page 255
- password (BGP), on page 257
- password-disable, on page 259
- password (rpki-server), on page 261
- permanent-network, on page 262

- precedence, on page 263
- preference (rpki-server), on page 265
- purge-time (rpki-server), on page 266
- rd, on page 267
- receive-buffer-size, on page 269
- redistribute (BGP), on page 271
- refresh-time (rpki-server), on page 276
- remote-as (BGP), on page 277
- remove-private-as, on page 280
- replace as-path all, on page 282
- response-time (rpki-server), on page 284
- retain local-label, on page 285
- retain route-target, on page 286
- route-monitoring inbound post-policy, on page 288
- route-monitoring local-rib, on page 289
- route-policy (BGP), on page 290
- route-reflector-client, on page 293
- router bgp, on page 295
- rpki route, on page 296
- rpki server, on page 298
- selective-vrf-download disable, on page 299
- send-buffer-size, on page 300
- send-community-ebgp, on page 302
- send-community-gshut-ebgp, on page 304
- send-extended-community-ebgp, on page 305
- session-group, on page 307
- session-open-mode, on page 309
- set flow-tag, on page 311
- show bgp, on page 312
- show bgp advertised, on page 327
- show bgp as-neighbors, on page 334
- show bgp af-group, on page 335
- show bgp aggregate-contributors, on page 338
- show bgp attribute-key, on page 340
- show bgp bmp, on page 344
- show bgp cidr-only, on page 370
- show bgp community, on page 374
- show bgp convergence, on page 379
- show bgp dampened-paths, on page 382
- show bgp flap-statistics, on page 386
- show bgp inconsistent-as, on page 391
- show bgp l2vpn, on page 395
- show bgp l2vpn vpls, on page 399
- show bgp labels, on page 405
- show bgp large-community, on page 408
- show bgp memory history, on page 411

- show bgp neighbor-group, on page 413
- show bgp neighbors, on page 417
- show bgp neighbors nsr, on page 436
- show bgp nexthops, on page 438
- show bgp nsr, on page 447
- show bgp paths, on page 451
- show bgp policy, on page 454
- show bgp process, on page 462
- show bgp regexp, on page 482
- show bgp route-policy, on page 486
- show bgp session-group, on page 490
- show bgp sessions, on page 493
- show bgp summary, on page 496
- show bgp summary nsr, on page 500
- show bgp table, on page 504
- show bgp truncated-communities, on page 507
- show bgp update-group, on page 511
- show bgp update in error process, on page 518
- show bgp update out, on page 519
- show bgp update out filter-group, on page 523
- show bgp update out process, on page 524
- show bgp update out sub-group, on page 526
- show bgp update out update-group, on page 528
- show bgp vrf, on page 530
- show bgp vrf update in error, on page 533
- show lpts punt excessive-flow-trap bgp, on page 534
- show protocols (BGP), on page 535
- show running-config lpts punt excessive-flow-trap, on page 537
- show svd role, on page 539
- show svd state, on page 540
- show tcp pcb, on page 541
- shutdown (BGP), on page 543
- shutdown (rpki-server), on page 545
- signalling disable, on page 546
- site-of-origin (BGP), on page 547
- slow peer (BGP neighbor address-family configuration), on page 549
- slow peer (BGP router global configuration), on page 552
- socket receive-buffer-size, on page 554
- socket send-buffer-size, on page 556
- soft-reconfiguration inbound, on page 558
- speaker-id, on page 561
- svd platform enable, on page 562
- table-policy, on page 564
- tcp mss inheritance-disable, on page 566
- timers (BGP), on page 567
- timers bgp, on page 569

- transport (rpki-server), on page 571
- ttl-security, on page 573
- update in error-handling basic disable, on page 576
- update in error-handling extended, on page 577
- update in error-handling treat-as-withdraw (BGP), on page 578
- update in labeled-unicast equivalent, on page 579
- update limit, on page 580
- update limit address-family, on page 581
- update limit sub-group, on page 583
- update out logging, on page 585
- update out originator-loopcheck disable, on page 586
- update-source, on page 588
- update wait-install, on page 590
- update wait-install delay startup, on page 591
- use, on page 592
- username (rpki-server), on page 597
- vrf (BGP), on page 598
- weight, on page 599

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 4.1.1	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

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

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#router bgp 100
RP/0/RSP0/CPU0:router(config-bgp)#neighbor 10.2.3.4
RP/0/RSP0/CPU0:router(config-bgp-nbr)#address-family vpnv4 unicast
RP/0/RSP0/CPU0:router(config-bgp-nbr-af)#accept-own
```

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 3.9.0	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

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

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#router bgp 100
RP/0/RSP0/CPU0:router(config-bgp)#address-family vpnv4 unicast
RP/0/RSP0/CPU0:router(config-bgp-af)#additional-paths install backup
```

Related Commands

Command	Description
retain local-label, on page 285	Retains the local label until the network is converged.

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 4.0.0	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.

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/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:routerconfig)# router bgp 100
RP/0/RSP0/CPU0:router(config-bgp)# address-family vpnv4 unicast
RP/0/RSP0/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/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config-bgp)#vrf vrf1
RP/0/RSP0/CPU0:router(config-bgp-vrf)#address-family ipv4 unicast
RP/0/RSP0/CPU0:router(config-bgp-vrf-af)#additional-paths receive disable
```

Related Commands

Command	Description
additional-paths send, on page 14	Configures send capability of multiple paths for a prefix to the capable peers.
capability additional-paths send, on page 128	Advertises capability of sending additional paths to the peer.
capability additional-paths receive, on page 126	Advertises additional paths receive capability.

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 4.0.0	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.

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 *Routing Configuration Guide for Cisco ASR 9000 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/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router bgp 100
RP/0/RSP0/CPU0:router(config-bgp)# address-family vpnv4 unicast
RP/0/RSP0/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/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config-bgp)#vrf vrf1
RP/0/RSP0/CPU0:router(config-bgp-vrf)#address-family ipv4 unicast
RP/0/RSP0/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/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config-bgp)#vrf vrf2
RP/0/RSP0/CPU0:router(config-bgp-vrf)#address-family ipv4 unicast
RP/0/RSP0/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 4.0.0	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.

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/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:routerconfig)# router bgp 100
RP/0/RSP0/CPU0:router(config-bgp)# address-family vpnv4 unicast
RP/0/RSP0/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/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config-bgp)#vrf vrf1
RP/0/RSP0/CPU0:router(config-bgp-vrf)#address-family ipv4 unicast
RP/0/RSP0/CPU0:router(config-bgp-vrf-af)#additional-paths send disable
```

Related Commands

Command	Description
additional-paths receive, on page 10	Configures receive capability of multiple paths for a prefix to the capable peers.
capability additional-paths send, on page 128	Advertises capability of sending additional paths to the peer.
capability additional-paths receive, on page 126	Advertises additional paths receive capability.

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** | **multicast** | **rt-filter** | **unicast** } | **ipv6** { **labeled-unicast** | **unicast** } | **l2vpn vpls-vpws** | **vpnv4** { **unicast**} | **vpnv6 {unicast }** }
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.
ipv6 unicast	Specifies IP Version 6 (IPv6) unicast address prefixes.
vpnv4 unicast	Specifies VPN Version 4 (VPNv4) unicast address prefixes. This option is not available in VRF or VRF neighbor configuration mode.
l2vpn vpls-vpws	Specifies L2VPN vpls-vpws address prefixes.
rt-filter	Specifies IPv4 RT-filter address 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

VRF configuration

VRF neighbor configuration (IPv4 address families)

Command History

Release	Modification
Release 3.9.0	L2VPN Address Family support was added.
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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

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.

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 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	—
l2vpn vpls-vpws	yes	yes	—
ipv4 rt-filter	yes	yes	—

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, IPv4 multicast, IPv4 labeled-unicast and VPNv4 unicast address families. IPv6 neighbor sessions support IPv6 unicast and IPv6 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:

```
Router# configure
Router(config)# router bgp 100
Router(config-bgp)# address-family ipv4 unicast
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:

```
Router# configure
Router(config)# router bgp 1
Router(config-bgp)# address-family ipv4 multicast
Router(config-bgp-af)# exit
Router(config-bgp)# neighbor 10.0.0.1
Router(config-bgp-nbr)# remote-as 1
Router(config-bgp-nbr)# address-family ipv4 multicast
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:

```
Router# configure
Router(config)# router bgp 12
Router(config-bgp)# address-family ipv4 tunnel
Router(config-bgp-af)#
```

advertise

To configure advertisement of local or re-originated VPNv4 or VPNv6 unicast routes or disable advertisement of L2VPN prefixes from a BGP router to its configured BGP neighbor, use the **advertise** command in BGP neighbor address family configuration mode. To undo this command configuration, use the **no** form of this command.

advertise { {**vpnv4** | **vpnv6**} **unicast** {**re-originated** | **local stitching-rt**} | **l2vpn evpn disable**}
no advertise { {**vpnv4** | **vpnv6**} **unicast** {**re-originated** | **local stitching-rt**} | **l2vpn evpn disable**}

Syntax Description

vpnv4	Specifies VPNv4 prefixes.
vpnv6	Specifies VPNv6 prefixes.
unicast	Specifies VPNv4 or VPNv6 unicast routes.
re-originated	Specifies advertisement of re-originated VPNv4 or VPNv6 unicast routes
local stitching-rt	Specifies advertisement of local VPNv4 or VPNv6 unicast routes with stitching route target identifier.
l2vpn	Specifies L2VPN address-family.
evpn disable	Disables advertisement of L2VPN EVPN prefixes.

Command Default

None

Command Modes

BGP neighbor address family configuration mode

Command History

Release	Modification
Release 5.3.2	This command was introduced.

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes 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 configure a BGP router to advertise local VPNv4 unicast routes assigned with stitching route target identifier to the specified BGP neighbor 1.1.1.1.

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router bgp 1
RP/0/RSP0/CPU0:router(config-bgp)# neighbor 1.1.1.1
RP/0/RSP0/CPU0:router(config-bgp-nbr)# address-family l2vpn evpn
RP/0/RSP0/CPU0:router(config-bgp-nbr-af)# advertise vpnv4 unicast re-originated stitching-rt
```

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 labelled unicast configuration

IPv6 labelled unicast configuration

Command History

Release	Modification
Release 3.9.0	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.

Unlabelled best-external is not supported as it may create routing loop.

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/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router bgp 100
RP/0/RSP0/CPU0:router(config-bgp)# address-family vpnv4 unicast
```

```
RP/0/RSP0/CPU0:router(config-bgp-af)# advertise best-external
```

Related Commands

Command	Description
additional-paths install backup, on page 9	Installs a backup path into the forwarding table and provides prefix independent convergence (PIC) in case of a PE-CE link failure.
retain local-label, on page 285	Retains the local label until the network is converged.

advertise local-labeled-route safi-unicast disable

To disable the advertisement of a route with a local label to peers on the unicast subaddress Family Identifier (SAFI), use the **advertise local-labeled-route safi-unicast disable** command in an appropriate address family configuration mode.

advertise local-labeled-route safi-unicast disable

Syntax Description

This command has no arguments or keywords.

Command Default

The default is to allow the path with local label to be advertised on the unicast SAFI.

Command Modes

IPv4 address family configuration

IPv6 address family configuration

Command History

Release	Modification
Release 6.2.1	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.

Example

To disable the advertisement of the route with a local label to peers on the unicast SAFI:

```
Router(config)# configure
Router(config)# router bgp 1
Router(config-bgp)# address-family ipv4 unicast
Router(config-bgp-af)# advertise local-labeled-route safi-unicast disable
```

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 5.1.1	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 shows how to advertise permanent path:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router bgp 100
RP/0/RSP0/CPU0:router(config-bgp)# neighbor 10.1.1.1
RP/0/RSP0/CPU0:router(config-bgp-nbr)# remote-as 4713
RP/0/RSP0/CPU0:router(config-bgp-nbr)# address-family ipv4 unicast
RP/0/RSP0/CPU0:router(config-bgp-nbr-af)# advertise permanent-network
```

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
```

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 3.7.2	This command was introduced.
Release 5.0.0	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.

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/RSP0/CPU0:router(config)# router bgp 5
RP/0/RSP0/CPU0:router(config-bgp)# neighbor 10.1.1.1
RP/0/RSP0/CPU0:router(config-bgp-nbr)# remote-as 100
RP/0/RSP0/CPU0:router(config-bgp-nbr)# advertisement-interval 10
```

Related Commands

Command	Description
neighbor-group, on page 231	Creates a neighbor group and enters neighbor group configuration mode.
session-group, on page 307	Creates a session group and enters session group configuration mode.

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 router configuration 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.
ipv6 unicast	Specifies IP Version 6 (IPv6) unicast address prefixes.
vpnv4 unicast	Specifies VPN Version 4 (VPNv4) unicast address prefixes.

Command Default

No BGP address family group is configured.

Command Modes

Router 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

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.

```
Router(config)# router bgp 100
Router(config-bgp)# af-group group1 address-family ipv4 unicast
Router(config-bgp-afgrp)# next-hop-self
```

Related Commands

Command	Description
neighbor (BGP), on page 229	Enters neighbor configuration mode for configuring BGP routing sessions.
neighbor-group, on page 231	Creates a neighbor group and enters neighbor group configuration mode.
session-group, on page 307	Creates a session group and enters session group configuration mode.
use, on page 592	Inherits configuration from a neighbor group, session group, or address family group.

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* | *aggregate-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.
route-policy *aggregate-route-policy-name*	(Optional) Specifies the name of a aggregate route policy.

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 3.7.2	This command was introduced.
Release 7.5.4	The command output displayed to set an aggregate contributor to a specific aggregate-address route.

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.

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.

Aggregate contributor flag is set in the router where you configure it. This feature is applicable for the following Address Family Indicators (AFIs):

- IPv4 unicast
- IPv6 unicast

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.

```
Router(config)# router bgp 100
Router(config-bgp)# address-family ipv4 unicast
Router(config-bgp-af)# aggregate-address 10.0.0.0/8 as-set
```

The following example shows to set an aggregate contributor to a specific aggregate-address route:

```
Router #config
Router(config)#router bgp 100
Router(config-bgp)#address-family ipv4 unicast
Router(config-bgp-af)#aggregate-address 250.2.2.0/24 route-policy aggregate-policy1
Router(config-bgp-af)#commit
```

Related Commands

Command	Description
network (BGP), on page 236	Specifies the list of networks for the BGP routing process.
route-policy (BGP), on page 290	Applies a routing policy to updates advertised to or received from a BGP neighbor
route-policy (RPL)	Defines a route policy and enters route-policy configuration mode.

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 4.0.0	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

Examples

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

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router bgp 100
RP/0/RSP0/CPU0:router(config-bgp)# neighbor 10.2.3.4
RP/0/RSP0/CPU0:router(config-bgp-nbr)# address-family ipv4 unicast
```

```
RP/0/RSP0/CPU0:router(config-bgp-nbr-af)# aigp
```

Related Commands

Command	Description
aigp send-cost-community	Sends AiGP value in cost community.

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 4.0.0	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.

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/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router bgp 100
RP/0/RSP0/CPU0:router(config-bgp)# neighbor 10.2.3.4
RP/0/RSP0/CPU0:router(config-bgp-nbr)# address-family ipv4 unicast
RP/0/RSP0/CPU0:router(config-bgp-nbr-af)# aigp send-cost-community 254
```

Related Commands

Command	Description
aigp, on page 33	Enables sending and receiving of accumulated interior gateway protocol (AiGP) attribute.

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 labeledunicast 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

VRF IPv4 address family 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

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 *MPLS Layer 3 VPN Configuration Guide for Cisco ASR 9000 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/RSP0/CPU0:router(config)# router bgp 6
RP/0/RSP0/CPU0:router(config-bgp)# address family ipv4 unicast
RP/0/RSP0/CPU0:router(config-bgp-af)# allocate-label route-policy policy_A
```

allow vpn default-originate

To configure the router to be enabled to advertise a default route to a configured BGP VPN neighbor, use the **allow vpn default-originate** command in the BGP VRF Address-Family configuration mode. To undo this configuration, use the **no** form of this command.

allow vpn default-originate
no allow vpn default-originate

Syntax Description

This command has no keywords or arguments.

Command Default

The router cannot advertise a default route to its BGP VPN neighbors.

Command Modes

BGP VRF Address-Family configuration mode

Command History

Release	Modification
Release 4.3.2	This command was introduced.

Usage Guidelines

This command only enables the router to advertise itself as the next-hop router for a default route to its BGP VPN neighbors. To actually forward the default route to a BGP VPN neighbor, you need to run the **default-originate** command under the BGP neighbor Address-Family configuration mode.

Task ID

Task ID	Operation
bgp	read, write

Example

The following example configuration shows how to enable a BGP router to advertise a default route to its BGP VPN neighbors.

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router bgp 1
RP/0/RSP0/CPU0:router(config-bgp)# vrf foo
RP/0/RSP0/CPU0:router(config-bgp-vrf)# address-family ipv4 unicast
RP/0/RSP0/CPU0:router(config-bgp-vrf-af)# allow vpn default-originate
```

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 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

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 ASR 9000 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/RSP0/CPU0:router(config)# router bgp 6
RP/0/RSP0/CPU0:router(config-bgp)# af-group group_1 address-family vpnv4 unicast
RP/0/RSP0/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 global configuration 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

Global configuration

Command History

Release	Modification
Release 3.9.0	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

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

```
RP/0/RSP0/CPU0:router(config)# as-format asdot
```

as-league peers

To configure a group of peer autonomous systems (AS) that will be used under common administration or a trusted relationship, use the **as-league peers** command in router configuration mode. To remove an AS from the group, use the **no** form of this command.

as-league peers *as-number*

Syntax Description

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

Command Default

None

Command Modes

Router configuration

Command History

Release	Modification
Release 4.2.1	Introduced in release 4.2.1.

Usage Guidelines

Additional paths send and receive capabilities are negotiated per neighbor per address family. Originally, Add-path capability was supported only for iBGP and confederation neighbors and not supported for eBGP neighbors. Use the **as-league peers** command to configure a subset of peer ASs so that the eBGP peers belonging to that league can participate in the add path send/receive. You can create a subset of peer ASs that are under a common administration, have a trusted relationship, or enable a select set of features among themselves.

Task ID

Task ID	Operation
bgp	read, write

This example shows how to create an AS league that includes a 200, 300, and 400. These are eBGP peers which, without an AS league peers configuration, do not negotiate additional paths send/receive. The following example shows how to create an AS league that includes peers 200, 300, and 400:

```
RP/0/RSP0/CPU0:router(config)# router bgp 100
RP/0/RSP0/CPU0:router(config-bgp)# as-league peers
RP/0/RSP0/CPU0:router(config-bgp-as-league-peers)# 200
RP/0/RSP0/CPU0:router(config-bgp-as-league-peers)# 300
RP/0/RSP0/CPU0:router(config-bgp-as-league-peers)# 400
RP/0/RSP0/CPU0:router(config-bgp-as-league-peers)#
```

!

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 3.7.2	This command was introduced.
Release 3.9.0	The **disable** keyword was replaced with the **inheritance-disable** keyword.
Release 7.11.1	This command was deprecated. The command remains available only to support OpenConfig configurations and is not intended for native CLI use.

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.

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/RSP0/CPU0:router(config)# router bgp 6
RP/0/RSP0/CPU0:router(config-bgp)# vrf vrf_A
RP/0/RSP0/CPU0:router(config-bgp-vrf)# neighbor 192.168.70.24
RP/0/RSP0/CPU0:router(config-bgp-vrf-nbr)# remote-as 10
RP/0/RSP0/CPU0:router(config-bgp-vrf-nbr)# address-family ipv4 unicast
RP/0/RSP0/CPU0:router(config-bgp-vrf-nbr-af)# as-override
```

Related Commands

Command	Description
site-of-origin (BGP), on page 547	Configures the site of origin filtering.

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 3.8.2	This command was introduced.
Release 7.11.1	This command was deprecated. The command remains available only to support OpenConfig configurations and is not intended for native CLI use.

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.

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/RSP0/CPU0:router#configure
```

```
RP/0/RSP0/CPU0:router(config)#router bgp 100
RP/0/RSP0/CPU0:router(config-bgp)#address-family ipv6 unicast
RP/0/RSP0/CPU0:router(config-bgp-af)#as-path-loopcheck out disable
```

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 4.2.3	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.

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:

```
Router#configure
Router(config)#router bgp 100
Router(config-bgp)#attribute-filter group ag_discard_med
Router(config-bgp-attrfg)#
```

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

```
Router#configure
Router(config)#router bgp 100
Router(config-bgp)#neighbor 10.0.1.101
Router(config-bgp-nbr)#remote-as 6461
Router(config-bgp-nbr)#update in filtering
Router(config-nbr-upd-filter)#attribute-filter group ag_discard_med
```

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

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

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/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#router bgp 3
RP/0/RSP0/CPU0:router(config-bgp)#attribute-filter group dogbert
RP/0/RSP0/CPU0:router(config-bgp-attrfg)#attribute large-community discard
RP/0/RSP0/CPU0:router(config-bgp-attrfg)#neighbor 10.10.10.3
RP/0/RSP0/CPU0:router(config-bgp-nbr)#update in filtering
RP/0/RSP0/CPU0:router(config-nbr-upd-filter)#attribute-filter group dogbert
```

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 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

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)

BFD value (global)	BFD value (local)	Result
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/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router bgp 65000
RP/0/RSP0/CPU0:router(config-bgp-nbrgrp)#neighbor 3.3.3.2
RP/0/RSP0/CPU0:router(config-bgp-nbr)# bfd minimum-interval 311
RP/0/RSP0/CPU0:router(config-bgp-nbr)# bfd multiplier 7
RP/0/RSP0/CPU0:router(config-bgp-nbr)# neighbor 5.5.5.2
RP/0/RSP0/CPU0:router(config-bgp-nbr)# bfd minimum-interval 318
RP/0/RSP0/CPU0:router(config-bgp-nbr)# bfd multiplier 4
RP/0/RSP0/CPU0:router(config-bgp-nbr)# vrf one
RP/0/RSP0/CPU0:router(config-bgp-vrf)# neighbor 3.12.1.2
RP/0/RSP0/CPU0:router(config-bgp-vrf-nbr)# bfd minimum-interval 119
RP/0/RSP0/CPU0:router(config-bgp-vrf-nbr)# bfd multiplier 10
RP/0/RSP0/CPU0:router(config-bgp-vrf-nbr)# commit
```

```
RP/0/RSP0/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/RSP0/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/RSP0/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/RSP0/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/RSP0/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
```

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 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 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

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

```
RP/0/RSP0/CPU0:router(config)# router bgp 6
RP/0/RSP0/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 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

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/RSP0/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/RSP0/CPU0:router(config)# router bgp 50
```

```
RP/0/RSP0/CPU0:router(config-bgp)# address-family ipv4 unicast
RP/0/RSP0/CPU0:router(config-bgp-af)# bgp attribute-download
!
!
!
RP/0/RSP0/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 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.



Note 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:

```
Router(config)# router bgp 6
Router(config-bgp)# bgp auto-policy-soft-reset disable
```

bgp bestpath aigp ignore

To configure a device that is running the Border Gateway Protocol (BGP) to not evaluate the accumulated interior gateway protocol (AIGP) metric during the best path selection process between two paths when one path does not have the AIGP metric, use the **bgp bestpath aigp ignore** command in router configuration mode. To return the device to default operation, use the **no** form of this command.

bgp bestpath aigp ignore
no bgp bestpath aigp ignore

Syntax Description

This command has no keywords or arguments.

Command Default

AIGP is enabled by default.

If this command is not configured, then the accumulated interior gateway protocol (AIGP) metric is evaluated (not ignored) during the best path selection.

Command Modes

Router configuration

VRF configuration

Command History

Release	Modification
Release 4.1.1	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.

By default, BGP always prefers a path with the AIGP metric. When there are two paths, one with the AIGP metric and the other without, then executing the **bgp bestpath aigp ignore** command results in BGP performing best path computation as if neither paths has the AIGP metric.

Task ID

Task ID	Operations
bgp	read, write

Examples

The following example shows how to configure the software to ignore the accumulated interior gateway protocol (AIGP) metric when performing best-path selection:

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#router bgp 50000
RP/0/RSP0/CPU0:router(config-bgp)#bgp bestpath aigp ignore
```

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 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

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/RSP0/CPU0:router(config)# router bgp 65000
RP/0/RSP0/CPU0:router(config-bgp)# bgp bestpath as-path ignore
```

Related Commands

Command	Description
bgp bestpath compare-routerid, on page 64	Compares identical routes received from eBGP peers during the best-path selection process and selects the route with the lowest router ID.

Command	Description
bgp bestpath med always, on page 67	Allows the comparison of the Multi Exit Discriminator (MED) for paths from neighbors in different autonomous systems.
bgp bestpath med confed, on page 69	Enables MED comparison among paths learned from confederation peers.
bgp bestpath med missing-as-worst, on page 71	Enables the software to consider a missing MED attribute in a path as having a value of infinity.

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 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

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/RSP0/CPU0:router(config)# router bgp 500
RP/0/RSP0/CPU0:router(config-bgp)# bgp bestpath compare-routerid
```

Related Commands

Command	Description
show bgp, on page 312	Displays entries in the BGP routing table.

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 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

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/RSP0/CPU0:router(config)# router bgp 500
RP/0/RSP0/CPU0:router(config-bgp)# bgp bestpath cost-community ignore
```

Related Commands

Command	Description
show bgp, on page 312	Displays entries in the BGP routing table.

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 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

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/RSP0/CPU0:router(config)# router bgp 100
RP/0/RSP0/CPU0:router(config-bgp)# bgp bestpath med always
```

Related Commands

Command	Description
bgp bestpath med confed, on page 69	Enables MED comparison among paths learned from confederation peers.
bgp bestpath med missing-as-worst, on page 71	Specifies that the software consider a missing MED attribute in a path as having a value of infinity, making the path without a MED value the least desirable path.
show bgp, on page 312	Displays entries in the BGP routing table.

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 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

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/RSP0/CPU0:router(config)# router bgp 210
RP/0/RSP0/CPU0:router(config-bgp)# bgp bestpath med confed
```

Related Commands

Command	Description
bgp bestpath med always, on page 67	Enables MED comparison among paths from neighbors in different autonomous systems.
bgp bestpath med missing-as-worst, on page 71	Specifies that the software consider a missing MED attribute in a path as having a value of infinity, making the path without a MED value the least desirable path.
show bgp, on page 312	Displays entries in the BGP routing table.

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 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 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

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/RSP0/CPU0:router(config)# router bgp 210
RP/0/RSP0/CPU0:router(config-bgp)# bgp bestpath med missing-as-worst
```

Related Commands

Command	Description
bgp bestpath med always, on page 67	Enables MED comparison among paths from neighbors in different autonomous systems.
bgp bestpath med confed, on page 69	Enables MED comparison among paths learned from confederation peers.

Command	Description
show bgp, on page 312	Displays entries in the BGP routing table.

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 global address family, neighbor, and neighbor address family submodes. To return the device to default operation, use the **no** form of this command.

For router and global address family configuration mode:

bgp bestpath origin-as allow invalid
no bgp bestpath origin-as allow invalid

For neighbour and neighbor address family configuration mode:

bestpath origin-as allow invalid
no 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

Neighbor

Neighbor Address family

Command History

Release	Modification
Release 4.2.1	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.

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.

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 in the global mode:

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#router bgp 50000
RP/0/RSP0/CPU0:router(config-bgp)#bgp bestpath origin-as allow invalid
```

The following example shows how to permit all invalid paths to be considered for BGP best-path selection in the address family submode:

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#router bgp 50000
RP/0/RSP0/CPU0:router(config-bgp)#address-family ipv4 unicast
RP/0/RSP0/CPU0:router(config-bgp-af)#bgp bestpath origin-as allow invalid
```

The following example shows how to permit all invalid paths to be considered for best-path selection in the neighbor submode:

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#router bgp 50000
RP/0/RSP0/CPU0:router(config-bgp)#neighbor 1.1.1.1
RP/0/RSP0/CPU0:router(config-bgp-nbr)#bestpath origin-as allow invalid
```

The following example shows how to permit all invalid paths to be considered for best-path selection in the neighbour address-family submode:

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#router bgp 50000
RP/0/RSP0/CPU0:router(config-bgp)#neighbor 1.1.1.1
RP/0/RSP0/CPU0:router(config-bgp-nbr)#address-family ipv4 unicast
RP/0/RSP0/CPU0:router(config-bgp-nbr-af)#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. To return the device to default operation, use the **no** form of this command.

bgp bestpath origin-as use validity
no 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 4.2.1	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.

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'.

The **bgp bestpath origin-as use validity** command takes precedence over the configured *weight* and *local preferences* attributes.

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:

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#router bgp 50000
RP/0/RSP0/CPU0:router(config-bgp)#bgp bestpath origin-as use validity
```

bgp bestpath as-path multipath-relax

To configure a Border Gateway Protocol (BGP) routing process to consider the different autonomous system (AS) paths and load balance multiple paths during best path route selection, use the **bgp bestpath as-path multipath-relax** command. To return the BGP routing process to the default operation, use the **no** form of this command.

bgp bestpath as-path multipath-relax
no bgp bestpath as-path multipath-relax

Syntax Description

This command has no keywords or arguments.

Command Default

None

Command Modes

Router BGP configuration

VRF 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

When BGP multi-pathing is enabled, BGP load-balances user traffic within a single autonomous system (AS). The criteria are that all attributes must match (weight, AS path, etc). However when a device is multi-homed to multiple autonomous systems, BGP cannot load balance traffic between them by default. In order to enable load-balancing of traffic among the multi-homed autonomous systems, the **bgp bestpath as-path multipath-relax** command needs to be enabled. The criteria required for this is that the AS-path length should be equal.

Before you use this command, ensure that BGP is enabled

Task ID

Task ID	Operation
bgp	read, write

Examples

This example shows how to configure multipath load sharing on paths from different autonomous systems in router mode:

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#router bgp 120
RP/0/RSP0/CPU0:router(config-bgp)#bgp bestpath as-path multipath-relax
```

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 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 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 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

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 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

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.

Examples

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

```
RP/0/RSP0/CPU0:router(config)# router bgp 65534
RP/0/RSP0/CPU0:router(config-bgp)# address-family ipv4 unicast
RP/0/RSP0/CPU0:router(config-bgp-af)# bgp client-to-client reflection disable
RP/0/RSP0/CPU0:router(config-bgp-af)# exit
RP/0/RSP0/CPU0:router(config-bgp)# neighbor-group rrclients
RP/0/RSP0/CPU0:router(config-bgp-nbrgrp)# remote-as 65534
RP/0/RSP0/CPU0:router(config-bgp-nbrgrp)# address-family ipv4 unicast
RP/0/RSP0/CPU0:router(config-bgp-nbrgrp-af)# route-reflector-client
RP/0/RSP0/CPU0:router(config-bgp-nbrgrp-af)# exit
RP/0/RSP0/CPU0:router(config-bgp-nbrgrp)# exit

RP/0/RSP0/CPU0:router(config-bgp)# neighbor 192.168.253.21 use neighbor-group rrclients
RP/0/RSP0/CPU0:router(config-bgp)# neighbor 192.168.253.22 use neighbor-group rrclients
```

Related Commands

Command	Description
bgp cluster-id, on page 80	Configures the cluster ID if the BGP cluster has more than one route reflector.
route-reflector-client, on page 293	Configures the router as a BGP route reflector and configures the specified neighbor as its client.

Command	Description
show bgp, on page 312	Displays entries in the BGP routing table.

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 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

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.

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/RSP0/CPU0:router(config)# router bgp 65534
RP/0/RSP0/CPU0:router(config-bgp)# bgp cluster-id 1
RP/0/RSP0/CPU0:router(config-bgp)# neighbor 192.168.70.24
RP/0/RSP0/CPU0:router(config-bgp-nbr)# remote-as 65534
RP/0/RSP0/CPU0:router(config-bgp-nbr)# address-family ipv4 unicast
RP/0/RSP0/CPU0:router(config-bgp-nbr-af)# route-reflector-client
```

Related Commands

Command	Description
route-reflector-client, on page 293	Configures the router as a BGP route reflector and configures the specified neighbor as its client.
show bgp, on page 312	Displays entries in the BGP routing table.

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 3.7.2	This command was introduced.
Release 3.9.0	Asplain format for 4-byte Autonomous system numbers notation was supported.

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.

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/RSP0/CPU0:router(config)# router bgp 4001
RP/0/RSP0/CPU0:router(config-bgp)# bgp confederation identifier 5
RP/0/RSP0/CPU0:router(config-bgp)# bgp confederation peers 4002
RP/0/RSP0/CPU0:router(config-bgp)# bgp confederation peers 4003
RP/0/RSP0/CPU0:router(config-bgp)# bgp confederation peers 4004
RP/0/RSP0/CPU0:router(config-bgp)# bgp confederation peers 4005
RP/0/RSP0/CPU0:router(config-bgp)# bgp confederation peers 4006
RP/0/RSP0/CPU0:router(config-bgp)# bgp confederation peers 4007
RP/0/RSP0/CPU0:router(config-bgp)# neighbor 10.2.3.4
RP/0/RSP0/CPU0:router(config-bgp-nbr)# remote-as 4002
RP/0/RSP0/CPU0:router(config-bgp-nbr)# exit
RP/0/RSP0/CPU0:router(config-bgp)# exit
RP/0/RSP0/CPU0:router(config-bgp-nbr)# neighbor 172.20.16.6
RP/0/RSP0/CPU0:router(config-bgp-nbr)# remote-as 4009
```

Related Commands

Command	Description
bgp confederation peers, on page 84	Configures the autonomous systems that belong to the confederation.

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 3.7.2	This command was introduced.
Release 3.9.0	Asplain format for 4-byte Autonomous system numbers notation was supported.

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.

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 82 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/RSP0/CPU0:router(config)# router bgp 1090
RP/0/RSP0/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/RSP0/CPU0:router(config)# router bgp 1095
RP/0/RSP0/CPU0:router(config-bgp)# bgp confederation peers
RP/0/RSP0/CPU0:router(config-bgp-confed-peers)# 1096
RP/0/RSP0/CPU0:router(config-bgp-confed-peers)# 1097
RP/0/RSP0/CPU0:router(config-bgp-confed-peers)# 1098
```

Related Commands

Command	Description
bgp confederation identifier, on page 82	Specifies a BGP confederation identifier.

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 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

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:

```
Router(config)# router bgp 50
Router(config-bgp)# address-family ipv4 unicast
Router(config-bgp-af)# bgp dampening 30 1500 10000 120
```

Related Commands

Command	Description
clear bgp dampening, on page 141	Clears BGP route dampening information and unsuppresses the suppressed routes.
clear bgp flap-statistics, on page 145	Clears BGP flap statistics.
route-policy (RPL)	Defines a route policy and enters route-policy configuration mode.
show bgp dampened-paths, on page 382	Displays BGP dampened routes.
show bgp flap-statistics, on page 386	Displays BGP flap statistics.
show bgp neighbors, on page 417	Displays information about BGP connections to neighbors.

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 3.9.0	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.

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:

```
Router(config)# router bgp 200
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 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

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/RSP0/CPU0:router(config)# router bgp 100
RP/0/RSP0/CPU0:router(config-bgp)# bgp enforce-first-as disable
```

Related Commands

Command	Description
show bgp, on page 312	Displays entries in the BGP routing table.

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 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

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/RSP0/CPU0:router(config)# router bgp 109
RP/0/RSP0/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 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

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/RSP0/CPU0:router(config)#router bgp 3
RP/0/RSP0/CPU0:router(config-bgp)#bgp graceful-restart
```

Related Commands

Command	Description
bgp graceful-restart graceful-reset, on page 93	Enables a graceful reset if configuration changes force a peer reset.
bgp graceful-restart purge-time, on page 94	Defines the maximum time before stale routes are purged.
bgp graceful-restart restart-time, on page 95	Defines the maximum time advertised to neighbors
bgp graceful-restart stalepath-time, on page 96	Defines the maximum time to wait for the End-of-RIB message from a neighbor that has been restarted before deleting learned routes.
show bgp, on page 312	Displays entries in the BGP routing table.
show bgp neighbors, on page 417	Displays information about BGP connections to neighbors.
show bgp process, on page 462	Displays BGP process information.

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 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

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/RSP0/CPU0:router(config)#router bgp 3
RP/0/RSP0/CPU0:router(config-bgp)# bgp graceful-restart graceful-reset
```

Related Commands

Command	Description
bgp graceful-restart, on page 91	Enables a graceful restart.
show bgp, on page 312	Displays entries in the BGP routing table.
show bgp neighbors, on page 417	Displays information about BGP connections to neighbors.
show bgp process, on page 462	Displays BGP process information.

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 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

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/RSP0/CPU0:router(config)# router bgp 3
RP/0/RSP0/CPU0:router(config-bgp)# bgp graceful-restart purge-time 800
```

Related Commands

Command	Description
bgp graceful-restart, on page 91	Enables a graceful restart.
show bgp, on page 312	Displays entries in the BGP routing table.
show bgp neighbors, on page 417	Displays information about BGP connections to neighbors.
show bgp process, on page 462	Displays BGP process information.

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 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

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/RSP0/CPU0:router(config)#router bgp 3
RP/0/RSP0/CPU0:router(config-bgp)# bgp graceful-restart restart-time 400
```

Related Commands

Command	Description
bgp graceful-restart, on page 91	Enables a graceful restart.
show bgp, on page 312	Displays entries in the BGP routing table.
show bgp neighbors, on page 417	Displays information about BGP connections to neighbors.
show bgp process, on page 462	Displays BGP process information.

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	Modifications
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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

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/RSP0/CPU0:router(config)# router bgp 3
RP/0/RSP0/CPU0:router(config-bgp)# bgp graceful-restart stalepath-time 750
```

Related Commands

Command	Description
bgp graceful-restart, on page 91	Enables a graceful restart.
show bgp, on page 312	Displays entries in the BGP routing table.
show bgp neighbors, on page 417	Displays information about BGP connections to neighbors.

Command	Description
show bgp process, on page 462	Displays BGP process information.

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 3.9.1	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

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

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#router bgp 100
RP/0/RSP0/CPU0:router(config-bgp)#address-family vpnv4 unicast
RP/0/RSP0/CPU0:router(config-bgp-af)#bgp import-delay 2 0
```

Related Commands

Command	Description
bgp label-delay, on page 99	Enables delay for Border Gateway Protocol (BGP) batch label processing

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 3.9.1	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

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

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#router bgp 100
RP/0/RSP0/CPU0:router(config-bgp)#address-family ipv4 unicast
```

```
RP/0/RSP0/CPU0:router(config-bgp-af)#bgp label-delay 2 0
```

Related Commands

Command	Description
bgp import-delay, on page 98	Enables delay for Border Gateway Protocol (BGP) batch import processing

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 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

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/RSP0/CPU0:router(config)# router bgp 65530
RP/0/RSP0/CPU0:router(config-bgp)# bgp log neighbor changes disable
```

Related Commands

Command	Description
show bgp neighbors, on page 417	Displays information about the TCP and BGP connections to neighbors.

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 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 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

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/RSP0/CPU0:router(config)#router bgp 65530
RP/0/RSP0/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 router configuration 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

Router configuration mode

Command History

Release	Modification
Release 5.2.0	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/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router bgp 100
RP/0/RSP0/CPU0:router(config-bgp)# bgp multipath as-path ignore onwards
```

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 6.5.1	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 IPv4 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
```

Examples

The following example shows how to enable origin-AS validation in IPv6 address-family configuration mode:

```
Router#configure
Router(config)#router bgp 50000
Router(config-bgp)#address-family ipv6 unicast
Router(config-bgp-af)#bgp origin-as validation enable
```

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 5.3.1	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 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

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/RSP0/CPU0:router(config)#router bgp 1
RP/0/RSP0/CPU0:router(config-bgp)# bgp redistribute-internal
RP/0/RSP0/CPU0:router(config-bgp)# exit
RP/0/RSP0/CPU0:router(config)# router ospf area1
RP/0/RSP0/CPU0:router(config-router)# redistribute bgp 1
RP/0/RSP0/CPU0:router(config-router)# end
```

```
RP/0/RSP0/CPU0:router# clear route *
```

Related Commands

Command	Description
clear bgp, on page 137 *	Resets all BGP neighbors.
clear route *	Resets all routes.

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 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

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 *Routing Configuration Guide for Cisco ASR 9000 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/RSP0/CPU0:router(config)# router bgp 100
RP/0/RSP0/CPU0:router(config-bgp)#bgp router-id 192.168.70.24
```

Related Commands

Command	Description
show bgp, on page 312	Displays entries in the BGP routing table.

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

VPNv6 address family 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

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/RSP0/CPU0:router(config)# router bgp 64500
RP/0/RSP0/CPU0:router(config-bgp)# address-family ipv4 unicast
RP/0/RSP0/CPU0:router(config-bgp-af)# bgp scan-time 20
```

Related Commands

Command	Description
show bgp, on page 312	Displays entries in the BGP routing table.

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 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

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/RSP0/CPU0:router(config)#router bgp 64530
RP/0/RSP0/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 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

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/RSP0/CPU0:router(config)# router bgp 65000
RP/0/RSP0/CPU0:router(config-bgp)#bgp write-limit 9000 27000
```

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

Release	Modification
6.6.3	This command was introduced.

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/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#router bgp 65000
RP/0/RSP0/CPU0:router(config-bgp)#rpki server 192.0.2.0
RP/0/RSP0/CPU0:router(config-bgp-rpki-server)# bind-source interface loopback 0
RP/0/RSP0/CPU0:router(config-bgp-rpki-server)#commit
RP/0/RSP0/CPU0:router(config-bgp-rpki-server)#no bind-source interface loopback 0
RP/0/RSP0/CPU0:router(config-bgp-rpki-server)#commit
```

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 5.2.2	This command was introduced.

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes 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

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/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router bgp 100
RP/0/RSP0/CPU0:router(config-bgp)# neighbor 1.1.1.1
RP/0/RSP0/CPU0:router(config-bgp-nbr)# bmp-activate server 4
```

Related Commands

Command	Description
bmp server, on page 123	Configures BMP server.
show bgp bmp, on page 344	Displays BMP information.

bmp advertisement-interval

To set the minimum interval between the sending of BMP routing updates, use the **advertisement-interval** command in router configuration mode.

Use the **no** form of this command to remove the **advertisement-interval** command from the configuration file and restore the system to its default interval values.

advertisement-interval *seconds*

Syntax Description

seconds	Minimum interval between sending BMP routing updates (in seconds). The range of the advertisement-interval of the route-monitoring inbound post-policy and the Local-RIB is from 2 seconds to 600 seconds. The default value is 15 seconds.

Command Default

None

Command Modes

Router configuration

Command History

Release	Modification
Release 7.5.4	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.
- You can configure **advertisement-interval** command for the following sub-modes to configure the interval between BMP route-monitoring updates:
  - **route-monitoring inbound post-policy**
  - **route-monitoring local-rib**

Task ID

Task ID	Operations
bgp	read, write

The following example configures an **advertisement-interval** for 15 seconds. This is to configure the interval between BMP route-monitoring updates for Inbound Post-Policy and Local-RIB:

```
Router#config
Router(config)#bmp server all
Router(config-bgp-bmp)#route-monitoring inbound post-policy
Router(config-bgp-bmp-rmon)#advertisement-interval 15
Router(config-bgp-bmp-rmon)#commit
```

```
Router#config
Router(config)#bmp server all
Router(config-bgp-bmp)#route-monitoring local-rib
Router(config-bgp-bmp-rmon)#advertisement-interval 15
Router(config-bgp-bmp-rmon)#commit
```

bmp scan-time

To configure scanning intervals of BMP-speaking networking devices, use the bmp **scan-time** command in router configuration mode.

Use the **no** form of this command to restore the scanning interval to its default value.

scan-time *seconds*

Syntax Description

seconds	Scanning interval (in seconds) of BMP routing information. The range of the scan time of route-monitoring inbound post-policy is from 5 seconds to 3600 seconds. The default value is 60 seconds.

Command Default

None

Command Modes

Router configuration

Command History

Release	Modification
Release 7.5.4	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

The following example displays the **scan-time** configuration for 60 seconds.

```
Router##config
Router#(config)#bmp server all
Router#(config-bgp-bmp)#route-monitoring inbound post-policy
Router#(config-bgp-bmp-rmon)#scan-time 60
Router#(config-bgp-bmp-rmon)#commit
```

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 Global Configuration mode. To remove a particular BMP server configuration, use **no** form of this command.

bmp server *server-id*

Syntax Description

server-id	Specifies BMP server ID. Server ID range is 1 to 8.
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

Global Configuration mode

Command History

Release	Modification
Release 5.2.2	This command was introduced.

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes 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

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/RSP0/CPU0:router# configure
RP/0/RSP0/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/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# bmp server 8 host 192.168.10.1 port 56
```

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

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# bmp server 5 update-source gigabitEthernet 0/0/0/1
```

Related Commands

Command	Description
bmp-activate, on page 119	Enables BMP logging for a neighbor.
show bgp bmp, on page 344	Displays BMP information.

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 4.0.0	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.

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/RSP0/CPU0:router(config)#router bgp 100
```

```
RP/0/RSP0/CPU0:router(config-bgp)#neighbor 10.2.3.4
RP/0/RSP0/CPU0:router(config-bgp-nbr)#capability additional-paths receive
```

Related Commands

Command	Description
additional-paths receive, on page 10	Configures receive capability of multiple paths for a prefix to the capable peers.
additional-paths send, on page 14	Configures send capability of multiple paths for a prefix to the capable peers.
capability additional-paths send, on page 128	Advertises capability of sending additional paths to the peer.

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 4.0.0	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.

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/RSP0/CPU0:router(config)# router bgp 100
```

```
RP/0/RSP0/CPU0:router(config-bgp)# neighbor 10.2.3.4
RP/0/RSP0/CPU0:router(config-bgp-nbr)# capability additional-paths send
```

Related Commands

Command	Description
additional-paths receive, on page 10	Configures receive capability of multiple paths for a prefix to the capable peers.
additional-paths send, on page 14	Configures send capability of multiple paths for a prefix to the capable peers.
capability additional-paths receive, on page 126	Advertises additional paths receive capability.

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 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

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/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# route-policy orfqq
RP/0/RSP0/CPU0:router:(config-rpl)# if orf prefix in (10.0.0.0/8 ge 20) then
RP/0/RSP0/CPU0:router(config-rpl)# pass
RP/0/RSP0/CPU0:router(config-rpl)# endif
RP/0/RSP0/CPU0:router(config-rpl)# if orf prefix in (1910::16 ge 120) then
RP/0/RSP0/CPU0:router(config-rpl)# pass
RP/0/RSP0/CPU0:router(config-rpl)# endif
RP/0/RSP0/CPU0:router(config-rpl)# end-policy
RP/0/RSP0/CPU0:router(config)# router bgp 65530
RP/0/RSP0/CPU0:router(config-bgp)# neighbor 10.0.101.1
RP/0/RSP0/CPU0:router(config-bgp-nbr)# remote-as 65534
RP/0/RSP0/CPU0:router(config-bgp-nbr)# address-family ipv4 unicast
RP/0/RSP0/CPU0:router(config-bgp-nbr-af)# route-policy pass-all out
RP/0/RSP0/CPU0:router(config-bgp-nbr-af)# capability orf prefix both
RP/0/RSP0/CPU0:router(config-bgp-nbr-af)# orf route-policy orfqq
```

Related Commands

Command	Description
af-group, on page 28	Creates an address family group for BGP neighbors and enters address family group configuration mode.
neighbor-group, on page 231	Creates a neighbor group and enters neighbor group configuration mode.
show bgp neighbors, on page 417	Displays information about BGP neighbors. Use the **received prefix-filter** keywords to display information on the prefix list filter.

capability suppress

To suppress capabilities from being advertised to the BGP peer, use the **capability suppress** command in the appropriate configuration mode. To remove the **capability suppress** 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** | **all** | **extended-nexthop-encoding**} [**inheritance-disable**]

Syntax Description

4-byte-as	Suppresses advertisement of the 4-byte AS capability to BGP peers.
all	Suppresses advertisement of all capabilities to BGP peers.
inheritance-disable	Prevents capability suppress 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 3.7.2	This command was introduced.
Release 3.9.0	The **disable** keyword was replaced with the **inheritance-disable** keyword.
Release 6.6.3	The **extended-nexthop-encoding** keyword 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/RSP0/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/RSP0/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/RSP0/CPU0:router(config)#router bgp 65000
RP/0/RSP0/CPU0:router(config-bgp)#neighbor 10.3.3.3
RP/0/RSP0/CPU0:router(config-bgp-nbr)#capability suppress 4-byte-as
RP/0/RSP0/CPU0:router(config-bgp-nbr)#commit
RP/0/RSP0/CPU0:router(config-bgp-nbr)#end

RP/0/RSP0/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/RSP0/CPU0:router(config-bgp)# neighbor 10.0.101.1
RP/0/RSP0/CPU0:router(config-bgp-nbr)# capability suppress 4-byte-as inheritance-disable

RP/0/RSP0/CPU0:router# show bgp neighbor 10.0.101.1 config
neighbor 10.0.101.1
 remote-as 1                 []
 address-family ipv4 unicast []

RP/0/RSP0/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 consistency-hashing auto-recovery

To enable automatic recovery of failed ECMP links and the sessions distributed due the ECMP link failure, use the **cef consistent-hashing auto-recovery** command in global configuration mode.

cef consistent-hashing auto-recovery

Syntax Description

This command has no keywords or arguments.

Command Default

Failed ECMP links are not automatically recovered.

Command Modes

Global configuration

Command History

Release	Modification
Release 6.5.1	The command was introduced.

Usage Guidelines

Configuring the command does not alter the current state. The command takes effect on the next link down or up events.

Task ID

Task ID	Operation
ipv4	read, write

Example

```
Router# configure
Router(config)# cef consistent-hashing auto-recovery
```

cef fast-reroute follow bgp-pic

A 4-minutes timer is being used to prevent traffic loss when using (Border Gateway Protocol) BGP with (Fast Reroute) FRR. The timer can cause an outage of up to 4 minutes when interface or BFD flaps are observed. The FRR would try to keep the backup route longer in use, but can cause a longer outage when interface or BFD flaps occurs. In order to avoid this potential outage created by the 4-minutes timer, the command **cef fast-reroute follow bgp-pic** is being used to disable it.

cef fast-reroute follow bgp-pic

Syntax Description

This command has no keywords or arguments.

Command Default

4-minutes timer acting upon when the FRR change is enabled.

Command Modes

Global configuration

Command History

Release	Modification
Release 6.3.2	The command was introduced.
Release 7.3.1	The command was deprecated.

Usage Guidelines

Configuring the command does not alter the current state. The command is used to disable the 4-minutes timer.

Task ID

Task ID	Operation
bgp	read, write

Example

```
Router# configure
Router(config)# cef fast-reroute follow bgp-pic
```

clear bgp

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

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

EXECEXEC mode

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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

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:

```
Router# clear bgp 10.0.0.1
```

Related Commands

Command	Description
clear bgp self-originated, on page 156	Clears self-originated routes.
clear bgp soft, on page 160	Soft resets a group of BGP neighbors.
show bgp, on page 312	Displays entries in the BGP routing table.
show bgp neighbors, on page 417	Displays information about the TCP and BGP connections to neighbors.

clear bgp current-mode

To switch from one BGP mode to another, use the **clear bgp current-mode** command in EXEC configuration mode.

clear bgp current-mode

Syntax Description This command has no arguments or keywords.

Command Default No default behavior or values

Command Modes EXEC 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **clear bgp current-mode** command to switch from standalone to distributed mode, or from distributed to standalone mode. The **show bgp process** command indicates the current BGP mode.

Note Switching from one mode to another causes all BGP sessions to go down.

Task ID

Task ID	Operations
bgp	execute

Examples The following example shows the **show bgp process** command output before and after switching from one BGP mode to another:

```
RP/0/RSP0/CPU0:router# show bgp process

  BGP Process Information
  BGP is operating in STANDALONE mode
  Autonomous System: 3
  Router ID: 10.18.18.11
  Cluster ID: 10.18.18.11
  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 not enabled
Client reflection is enabled
Scan interval: 60
Main Table Version: 3
IGP notification: IGPs notified
RIB has converged: version 0

Node                 Process     Nbrs Estb Rst Upd-Rcvd Upd-Sent Nfn-Rcv Nfn-Snt
node0_0_CPU0         Speaker        5    5  51        0        7       0       5

RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router bgp 3
RP/0/RSP0/CPU0:router(config-bgp)# distributed speaker 1
RP/0/RSP0/CPU0:router(config-bgp)# distributed speaker 2
RP/0/RSP0/CPU0:router(config-bgp)# neighbor 10.0.101.1
RP/0/RSP0/CPU0:router(config-bgp-nbr)# speaker-id 2
RP/0/RSP0/CPU0:router(config-bgp-nbr)# commit
RP/0/RSP0/CPU0:router# clear bgp current-mode
RP/0/RSP0/CPU0:router# show bgp process

BGP Process Information
BGP is operating in DISTRIBUTED mode
Autonomous System: 3
Router ID: 10.18.18.11
Cluster ID: 10.18.18.11
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 not enabled
Client reflection is enabled
Scan interval: 60
Main Table Version: 1
IGP notification: IGPs not notified
RIB has not converged: version 0

Node                 Process     Nbrs Estb Rst Upd-Rcvd Upd-Sent Nfn-Rcv Nfn-Snt
node0_0_CPU0         Speaker 1      4    1  52        0        0       0       4
node0_0_CPU0         Speaker 2      1    0   9        0        0       0       1
node0_0_CPU0         bRIB 1         0    0   0        0        0       0       0
node0_0_CPU0         bRIB 2         0    0   0        0        0       0       0
```

Related Commands

Command	Description
show bgp process, on page 462	Displays the current BGP process information.

clear bgp dampening

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

clear bgp dampening

Syntax Description

ipv4	Specifies IP Version 4 address prefixes.
unicast	Specifies unicast address prefixes.
multicast	Specifies multicast 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.
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

EXEC configurationEXEC mode

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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

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:

```
Router# clear bgp ipv4 unicast dampening 172.20.0.0/16
```

Related Commands

Command	Description
bgp dampening, on page 86	Enables BGP route dampening or changes various BGP route dampening factors.
show bgp dampened-paths, on page 382	Displays BGP dampened routes.

clear bgp external

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

EXEC 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 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	execute

Examples

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

```
RP/0/RSP0/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 EXEC configurationEXEC mode 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

EXEC

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 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	execute

Examples

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

```
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:

```
Router# clear bgp ipv4 unicast flap-statistics 172.20.1.1
```

clear bgp long-lived-stale

To delete all paths received from the given neighbor that are long-lived-stale, use the **clear bgp long-lived-stale** command in EXEC mode.

clear bgp vrf {*vrf-name* | **all**} {**ipv4** | **ipv6**} **unicast** *nbr-address* **long-lived-stale**

Syntax Description

vrf *vrf-name*	Deletes all paths received from the given neighbor that are long-lived-stale for the specified VRF
vrf all	Deletes all paths received from the given neighbor that are long-lived-stale for all VRFs.
ipv4 unicast	Specifies IP Version 4 unicast address prefixes.
ipv6 unicast	Specifies IP Version 6 unicast address prefixes.
nbr-address	Specifies IPv4 or IPv6 address of the neighbor.

Command Default

No default behavior

Command Modes

EXEC

Command History

Release	Modification
Release 5.2.2	This command was introduced

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes 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

This command deletes all paths received from the given neighbor for all VRFs:

```
RP/0/0/CPU0:router# clear bgp vrf all ipv4 unicast 192.172.20.10 long-lived-stale
```

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 EXEC configuration 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.
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.

Command Default

No default behavior or values

Command Modes

EXEC 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

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/RSP0/CPU0:router# clear bgp vrf vrf_A nexthop performance statistics
```

Related Commands

Command	Description
show bgp nexthops, on page 438	Displays information about the BGP next-hop notifications.

clear bgp nexthop registration

To reregister a specified next hop with the Routing Information Base (RIB), use the **clear bgp nexthop registration** command in EXEC configuration 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.
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.
nexthop-address	Address of the next hop.

Command Default

No default behavior or values

Command Modes

EXEC 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

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/RSP0/CPU0:router# clear bgp nexthop registration 10.1.1.1
```

Related Commands

Command	Description
show bgp nexthops, on page 438	Displays information about the BGP next-hop notifications.

clear bgp peer-drops

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

EXEC 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 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	execute

Examples

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

```
RP/0/RSP0/CPU0:router# clear bgp peer-drops *
```

Related Commands

Command	Description
show bgp neighbors, on page 417	Displays information about BGP connections to neighbors.

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.

Command Default

No default behavior or values

Command Modes

EXEC 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 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	execute

Examples

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

```
RP/0/RSP0/CPU0:router# clear bgp performance-statistics
```

clear bgp rpki

To clear resource public key infrastructure (RPKI) cache-server (rpki-sever) or Prefix validation state, use the `clear bgp rpki` command.

Use the `server` option in the `clear bgp rpki` command in EXEC mode to clear the Resource Public Key Infrastructure (RPKI) server, purge Route Origin Authorizations (ROA) records downloaded from that server, renegotiate the TCP connection, and redownload ROA records.

Use `clear bgp rpki server WORD` command to clear a specific RPKI cache-server (hostname or IP address).

The `clear bgp rpki server *` command clears all RPKI cache-servers.

Use the `clear bgp rpki validation` command to re-run validation for all prefixes under the specified VRF or all VRFs, if not specified. If the address family is specified after the command, then only prefixes belonging to that address family will have their validation checks re-run.

clear bgp rpki server { ***** | **WORD** }

clear bgp { **vrf <vrf name>** } **rpki validation** { **ipv4** | **ipv6** }

Syntax Description

server	Clear RPKI cache-server.
validation	Clear RPKI origin-validation (trigger origin validation).
*****	Clear all RPKI cache-servers.
WORD	Clear a specific RPKI cache-server (hostname or IP address).
ipv4	Trigger origin validation walk for the IPv4 table.
ipv6	Trigger origin validation walk for the IPv6 table.

Command Modes

EXEC

Command History

Release	Modification
Release 4.2.1	This command was introduced.

The following example shows how to configure BGP RPKI server:

```
router bgp 65000
 rpki server 10.0.101.2
  bind-source interface GigabitEthernet0/2/0/1
  transport tcp port 32001
  response-time 15
```

The following example shows the output of `show bgp rpki server summary` and `show bgp rpki table` commands:

```
RP/0/0/CPU0:router#show bgp rpki server summary
Wed May  6 10:26:42.532 PDT
```

```
Hostname/Address        Transport       State           Time            ROAs (IPv4/IPv6)
10.0.101.2              TCP:32001       ESTAB           00:00:09        3/3

RP/0/0/CPU0:router#show bgp rpki table
Wed May  6 10:26:44.592 PDT

  Network               Maxlen          Origin-AS         Server
  1.1.1.0/24            32              9991              10.0.101.2
  67.64.0.0/30          32              4002              10.0.101.2
  67.128.0.0/30         32              4003              10.0.101.2

  Processed 3 RPKI entries
```

The following example shows the output of `show bgp rpki server summary` command after clearing RPKI server 10.0.101.2:

```
RP/0/0/CPU0:router#clear bgp rpki server 10.0.101.2
Wed May  6 10:26:50.318 PDT

RP/0/0/CPU0:router#show bgp rpki server summary
Wed May  6 10:26:53.212 PDT

Hostname/Address        Transport       State           Time            ROAs (IPv4/IPv6)
10.0.101.2              TCP:32001       NONE            00:00:04        N/A

RP/0/0/CPU0:router#show bgp rpki table
Wed May  6 10:26:55.837 PDT

  Processed 0 RPKI entries

RP/0/0/CPU0:router#show bgp rpki server summary
Wed May  6 10:27:04.663 PDT

Hostname/Address        Transport       State           Time            ROAs (IPv4/IPv6)
10.0.101.2              TCP:32001       ESTAB           00:00:11        0/0

RP/0/0/CPU0:router#show bgp rpki table
Wed May  6 10:27:17.513 PDT

  Network               Maxlen          Origin-AS         Server
  1.1.1.0/24            32              9991              10.0.101.2
  67.64.0.0/30          32              4002              10.0.101.2
  67.128.0.0/30         32              4003              10.0.101.2

  Processed 3 RPKI entries
```

clear bgp self-originated

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

Command Default

No default behavior or values

Command Modes

EXEC configurationEXEC mode

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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

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:

```
Router# clear bgp ipv4 unicast self-originated
```

Related Commands

Command	Description
aggregate-address, on page 30	Creates an aggregate entry in a BGP routing table.
network (BGP), on page 236	Specifies a local network that the BGP routing process should originate and advertise to its neighbors.
redistribute (BGP), on page 271	Redistributes routes from another routing protocol into BGP.

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 EXEC configuration 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.

Command Default

No default behavior or values

Command Modes

EXEC 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 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	execute

Examples

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

```
RP/0/RSP0/CPU0:router# clear bgp shutdown
```

Related Commands

Command	Description
show bgp, on page 312	Displays entries in the BGP routing table.
show bgp neighbors, on page 417	Displays information about the TCP and BGP connections to neighbors.

clear bgp soft

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

EXEC configurationEXEC mode

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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

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
```

Related Commands

Command	Description
bgp auto-policy-soft-reset disable, on page 60	Disables an automatic soft reset of BGP peers when the configured inbound route policy is modified.
clear bgp, on page 137	Resets a group of BGP neighbors.
clear bgp self-originated, on page 156	Clears self-originated routes.
show bgp, on page 312	Displays entries in the BGP routing table.
show bgp neighbors, on page 417	Displays information about the TCP and BGP connections to neighbors.
soft-reconfiguration inbound, on page 558	Configures the software to store updates received from a neighbor.

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 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

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/RSP0/CPU0:router(config)#router bgp 164
RP/0/RSP0/CPU0:router(config-bgp)# default-information originate
```

Related Commands

Command	Description
redistribute (BGP), on page 271	Redistributes routes from another protocol into BGP.

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

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router bgp 100
RP/0/RSP0/CPU0:router(config-bgp)# address-family ipv4 multicast
RP/0/RSP0/CPU0:router(config-bgp-af)# default-martian-check disable
```

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

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router bgp 100
RP/0/RSP0/CPU0:router(config-bgp)# address-family ipv6 multicast
RP/0/RSP0/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 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

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/RSP0/CPU0:router(config)# router bgp 109
RP/0/RSP0/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

VRF IPv4 neighbor address family configuration

VRF IPv6 neighbor address family configuration

Command History

Release	Modification
Release 3.7.2	This command was introduced.
Release 3.9.0	The **disable** keyword was replaced with the **inheritance-disable** keyword.

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.

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.

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/RSP0/CPU0:router(config)# router bgp 109
RP/0/RSP0/CPU0:router(config-bgp)#address-family ipv4 unicast
RP/0/RSP0/CPU0:router(config-bgp)# neighbor 172.20.2.3
RP/0/RSP0/CPU0:router(config-bgp-nbr)# remote-as 200
RP/0/RSP0/CPU0:router(config-bgp-nbr)# address-family ipv4 unicast
RP/0/RSP0/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/RSP0/CPU0:router(config)# router bgp 109
RP/0/RSP0/CPU0:router(config-bgp)# neighbor 172.20.2.3
RP/0/RSP0/CPU0:router(config-bgp-nbr)# remote-as 200
RP/0/RSP0/CPU0:router(config-bgp-nbr)# address-family ipv4 unicast
RP/0/RSP0/CPU0:router(config-bgp-nbr-af)# default-originate route-policy
default-default-policy
```

Related Commands

Command	Description
default-information originate (BGP), on page 164	Allows the default route to be redistributed into BGP from another routing protocol.
af-group, on page 28	Creates an address family group for BGP neighbors and enters address family group configuration mode.
neighbor-group, on page 231	Creates a neighbor group and enters neighbor group configuration mode.

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 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

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/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)#router bgp 65000
RP/0/RSP0/CPU0:router(config-bgp)#neighbor 192.168.13.4
RP/0/RSP0/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 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

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/RSP0/CPU0:router(config)# router bgp 109
RP/0/RSP0/CPU0:router(config-bgp)# address-family ipv4 unicast
RP/0/RSP0/CPU0:router(config-bgp-af)#distance bgp 20 20 200
```

Related Commands

Command	Description
distance (IS-IS)	Defines the administrative distance assigned to routes discovered by the IS-IS protocol.
distance (OSPF)	Defines OSPF route administrative distances based on route type.

distribute bgp-ls (ISIS)

To distribute ISIS link-state data using BGP LS, use the **distribute bgp-ls** command in router configuration 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

Router configuration.

Command History

Release	Modification
Release 5.1.1	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
isis	read, write

Examples

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

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router isis foo
RP/0/RSP0/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 router configuration 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

Router configuration.

Command History

Release	Modification
Release 5.1.1	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
ospf	read, write

Examples

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

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router ospf 100
RP/0/RSP0/CPU0:router(config-ospf)# distribute bgp-ls instance-id 32 throttle 10
```

distributed speaker

To enable a distributed speaker process, use the **distributed speaker** command in router configuration mode. To remove the distributed speaker process, use the **no** form of this command.

distributed speaker *id*
no distributed speaker *id*

Syntax Description

id	ID of the distributed speaker process. Range is 1 to 15.

Command Default

Default is 0.

Command Modes

Router 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

If BGP is currently running in standalone mode, you must enter the **clear bgp current-mode** command to switch from standalone or distributed mode.

Task ID

Task ID	Operations
bgp	read, write

Examples

The following example shows how to enable distributed speaker process 3:

```
RP/0/RSP0/CPU0:router(config)# router bgp 109
RP/0/RSP0/CPU0:router(config-bgp)# distributed speaker 3
```

Related Commands

Command	Description
clear bgp current-mode, on page 139	Switch BGP from one mode to another

dmz-link-bandwidth

To originate a demilitarized zone (DMZ) link bandwidth extended community for the link to an eBGP or iBGP 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

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.

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:

```
Router(config)# router bgp 1
Router(config-bgp)#neighbor 45.67.89.01
Router(config-bgp-nbr)#dmz-link-bandwidth
```

Related Commands

Command	Description
bandwidth	Configures the bandwidth of an interface.
maximum-paths (BGP)	Controls the maximum number of parallel routes that Border Gateway Protocol (BGP) installs in the routing table.

Command	Description
neighbor-group	Creates a neighbor group and enters neighbor group configuration mode.
session-group	Creates a session group and enters session group configuration mode.

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 5.1.1	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 shows how to configure a unique identifier ASN:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router bgp 100
RP/0/RSP0/CPU0:router(config-bgp)# address-family link-state link-state
RP/0/RSP0/CPU0:router(config-bgp-af)# domain-distinguisher 1234
```

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 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

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

Table 2: dscp Default Settings, on page 178 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

DSCP (Precedence)	Minimum Threshold	Maximum Threshold	Mark Probability
af23	24	40	1/10
af31	32	40	1/10
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/RSP0/CPU0:router(config)# router bgp 5
RP/0/RSP0/CPU0:router(config-bgp)#neighbor 10.1.1.1
RP/0/RSP0/CPU0:router(config-bgp-nbr)#remote-as 100
RP/0/RSP0/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 3.7.2	This command was introduced.
Release 4.0.0	The **mpls** keyword was supported.

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.

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/RSP0/CPU0:router(config)# router bgp 109
RP/0/RSP0/CPU0:router(config-bgp)# neighbor 172.20.16.6
RP/0/RSP0/CPU0:router(config-bgp-nbr)# ebgp-multihop
```

Related Commands

Command	Description
neighbor-group, on page 231	Creates a neighbor group and enters neighbor group configuration mode.
session-group, on page 307	Creates a session group and enters session group configuration mode.

ecmp-delay

To configure delay in best path calculation for prefixes with ECMP paths in BGP, use the **ecmp-delay** submode in the address-family configuration mode.

ecmp-delay { **fixed delay** *delay_interval_value* | **as-based delay** *delay_interval_value* [**as-list** *as_list_name*] | **platform-oor-based delay** *delay_interval_value* **threshold** *threshold_value* }

Syntax Description

as-based	Enables ECMP delay based on the neighbor AS.
as-list *as_list_name*	Enables ECMP delay for neighbor AS included in the AS list.
fixed	Enables ECMP delay with a fixed delay.
platform-oor-based	Enables ECMP delay with an OOR threshold.
delay	Specifies the delay interval in milliseconds.
delay_interval_value	Specifies the duration that the router delays the RIB or FIB installation for BGP prefixes. The range is 10–300000.
threshold *threshold_value*	Applies ECMP delay with an OOR threshold value, only when the platform resource usage is above the specified threshold value.

Command Default

No default behavior or values

Command Modes

IPv4 or IPv6 address-family configuration

Command History

Release	Modification
Release 25.1.1	This command was introduced.

Usage Guidelines

None.

Task ID

Task ID	Operations
bgp	read, write

Examples

The following example shows how to configure an ECMP delay of 10 milliseconds based on the neighbor AS:

```
Router(config)#router bgp 100
Router(config-bgp)#address-family ipv4 unicast
Router(config-bgp-af)#maximum-paths eibgp 1024 selective route-policy mp_rpl
Router(config-bgp-af)#ecmp-delay
Router(config-bgp-af-ecmpdelay)#as-based delay 10
Router(config-bgp-af)#commit
```

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 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

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/RSP0/CPU0:router(config)# vrf vrf-1
RP/0/RSP0/CPU0:router(config-vrf)#address-family ipv4 unicast
RP/0/RSP0/CPU0:router(config-vrf-af)# export route-policy policy-A
```

Related Commands

Command	Description
import route-policy, on page 195	Specifies a route policy to import routes into the VRF instance.

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 3.7.2	This command was introduced.
Release 3.9.0	Asplain format for 4-byte Autonomous system numbers notation was supported.

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.

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
ip-services	read, write

Examples

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

```
RP/0/RSP0/CPU0:router(config)# vrf vrf-1
RP/0/RSP0/CPU0:router(config-vrf)# address-family ipv4 unicast
RP/0/RSP0/CPU0:router(config-vrf-af)# export route-target 500:1
```

Related Commands

Command	Description
import route-target, on page 196	Specifies the import route-target.

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 5.3.2	This command was introduced.

Task ID

Task ID	Operations
bgp	read, write

host-reachability protocol bgp

To configure a VxLAN Tunnel EndPoint (VTEP) with BGP as the control plane that provides BGP protocol-based VTEP peer discovery and end-host reachability information distribution, use the **host-reachability protocol bgp** command in the VNI NVE configuration mode. To revert to the default configuration, use the no form of this command.

host-reachability protocol bgp
no host-reachability protocol bgp

Syntax Description

This command has no keywords or arguments.

Command Default

The VTEP relies on a multicast-based data plane flood-and-learn behavior for VTEP peer discovery and end-host reachability information as defined in the initial IETF VXLAN standards (RFC 7348).

Command Modes

VNI NVE configuration mode

Command History

Release	Modification
Release 5.3.2	This command was introduced.

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes 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
tunnel, interface	read, write

Example

The following example configuration shows BGP configured as the control plane for a Network Virtualization EndPoint (NVE) interface (VTEP) that is associated with a VxLAN with identifier 1.

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# interface nve 1
RP/0/RSP0/CPU0:router(config-nve)# member vni 1
RP/0/RSP0/CPU0:router(config-nve-vni)# host-reachability protocol bgp
```

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 router configuration 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

Router configuration

Command History

Release	Modification
Release 7.8.1	The command scope was expanded to influence next-hop-self behavior, enforce outbound policy modifications, and optimize label processing in BGP. Refer to the Usage Guidelines for more information.
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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

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.

For VPN or EVPN with Route Reflectors (RRs) configured with a next-hop-self, no label allocation occurs unless the **ibgp policy out enforce-modifications** option is configured.

In Option B scenarios with the same Route Distinguisher (RD) and using RRs with next-hop-self, for example, ABR, label allocation does not occur unless the **ibgp policy out enforce-modifications** option is configured.

RIB (Routing Information Base) processing adheres to the label versioning. Labels are processed before RIB processing and then considered for updates. If an RR is configured with the **ibgp policy out enforce-modifications** option, label processing will occur; otherwise, it will be optimized to update the RIB and process updates without label processing.

Outbound policy modifications on RRs, such as modifying extcomm and largecomm, are only applied when the **ibgp policy out enforce-modifications** option is configured.

Modifying the AS Path using Prepend and Replace-AS in RRs is effective only with the **ibgp policy out enforce-modifications** configuration.

For RRs with multipath or next-hop-self, the selected label during sending updates is the local label only when configured with the **ibgp policy out enforce-modifications** configuration.

Next-hop-self in RRs is applied with the **ibgp policy out enforce-modifications** configuration.

For 6PE detection where a labeled path is from a V4 Neighbor, implicit next-hop is applied when advertising an IPv6 Unicast. Cross AFI (Address Family Identifier) labeled unicast is not advertised to IPv6 unicast if the **ibgp policy out enforce-modifications** option is not configured.

Task ID

Task ID	Operations
bgp	read, write

Examples

The following example shows how to set the ibgp policy out enforce-modifications:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router bgp 6500
RP/0/RSP0/CPU0:router(config-bgp)# ibgp policy out enforce-modifications
```

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 3.9.0	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

The following example shows how to enable ignore-connected check configuration for neighbor 10.2.3.4:

```
RP/0/RSP0/CPU0:router(config)# router bgp 100
RP/0/RSP0/CPU0:router(config-bgp)# neighbor 10.2.3.4
RP/0/RSP0/CPU0:router(config-bgp-nbr)# ignore-connected-check
```

import

To configure the import and export of BGP NLRIs between two BGP neighbors with respect to the route target identifiers, use the **import** command in the BGP neighbor address family configuration mode. To undo this command configuration, run the **no** form of this command.

import [**stitching-rt**] **reoriginate** [**stitching-rt**]

no import [**stitching-rt**] **reoriginate** [**stitching-rt**]

Syntax Description

reoriginate	Imports NLRIs that match normal route target identifier and exports re-originated NLRIs assigned with the normal route target identifier.
reoriginate stitching-rt	Imports NLRIs that match normal route target identifier and exports re-originated NLRIs assigned with the stitching route target identifier
stitching-rt reoriginate	Imports NLRIs that match stitching route target identifier and exports re-originated NLRIs assigned with the normal route target identifier.
stitching-rt reoriginate stitching-rt	Imports NLRIs that match stitching route target identifier and exports re-originated NLRIs assigned with the stitching route target identifier.

Command Default

None

Command Modes

BGP neighbour address family configuration mode

Command History

Release	Modification
Release 5.3.2	This command was introduced.

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes 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 configure DCI router to import BGP NLRIs that match normal route target identifier and to export re-originated BGP NLRIs assigned with the stitching route target identifier.

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router bgp 1
RP/0/RSP0/CPU0:router(config-bgp)# neighbor 1.1.1.1
RP/0/RSP0/CPU0:router(config-bgp-nbr)# address-family l2vpn evpn
RP/0/RSP0/CPU0:router(config-bgp-nbr-af)# import reoriginate stitching-rt
```

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 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

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/RSP0/CPU0:router(config)# vrf vrf-1
RP/0/RSP0/CPU0:router(config-vrf)# address-family ipv4 unicast
RP/0/RSP0/CPU0:router(config-vrf-af)# import route-policy policy-B
```

Related Commands

Command	Description
export route-policy, on page 184	Specifies a route policy to export routes from the VRF instance.

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 3.7.2	This command was introduced.
Release 3.9.0	Asplain format for 4-byte Autonomous system numbers notation was supported.

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.

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/RSP0/CPU0:router(config)#vrf vrf-1
RP/0/RSP0/CPU0:router(config-vrf)# address-family ipv4 unicast
RP/0/RSP0/CPU0:router(config-vrf-af)# import route-target 500:99
```

Related Commands

Command	Description
export route-target, on page 185	Specifies the export route-target.

keychain

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 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

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/RSP0/CPU0:router(config)# router bgp 140
RP/0/RSP0/CPU0:router(config-bgp)# neighbor 172.20.1.1
RP/0/RSP0/CPU0:router(config-bgp-nbr)# remote-as 1
RP/0/RSP0/CPU0:router(config-bgp-nbr)# keychain keychain_A
```

Related Commands

Command	Description
keychain-disable, on page 200	Overrides any inherited key chain configuration from a neighbor group or session group for BGP neighbors.

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 202 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 3.7.2	This command was introduced.
Release 3.9.0	This command was replaced by the **keychain inheritance-disable** command.

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.

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/RSP0/CPU0:router(config)# router bgp 140
RP/0/RSP0/CPU0:router(config-bgp)# session-group group1
RP/0/RSP0/CPU0:router(config-bgp-sngrp)# keychain keychain_A
RP/0/RSP0/CPU0:router(config-bgp-sngrp)# exit
RP/0/RSP0/CPU0:router(config-bgp)#neighbor 172.20.1.1
RP/0/RSP0/CPU0:router(config-bgp-nbr)# remote-as 2
RP/0/RSP0/CPU0:router(config-bgp-nbr)#use session-group group1
RP/0/RSP0/CPU0:router(config-bgp-nbr)# keychain-disable
```

Related Commands

Command	Description
keychain, on page 198	Enables key chain authentication on a TCP connection between two BGP neighbors.

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 3.9.0	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.

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/RSP0/CPU0:router(config)#router bgp 140
RP/0/RSP0/CPU0:router(config-bgp)# session-group group1
RP/0/RSP0/CPU0:router(config-bgp-sngrp)# keychain keychain_A
RP/0/RSP0/CPU0:router(config-bgp-sngrp)# exit
RP/0/RSP0/CPU0:router(config-bgp)# neighbor 172.20.1.1
RP/0/RSP0/CPU0:router(config-bgp-nbr)# remote-as 2
RP/0/RSP0/CPU0:router(config-bgp-nbr)# use session-group group1
RP/0/RSP0/CPU0:router(config-bgp-nbr)# keychain inheritance-disable
```

Related Commands

Command	Description
keychain, on page 198	Enables key chain authentication on a TCP connection between two BGP neighbors.

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**]
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.

Command Default

Per-prefix is the default label allocation mode.

Command Modes

VRF configuration

Command History

Release	Modification
Release 3.7.2	This command was introduced.
Release 7.4.1	This command was deprecated

Note The function of this deprecated command can be carried out using label mode command under configured address-family.

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.

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.

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/RSP0/CPU0:router(config)# router bgp 109
```

```
RP/0/RSP0/CPU0:router(config-bgp)# vrf vrf-1
RP/0/RSP0/CPU0:router(config-bgp-vrf)# label-allocation-mode per-ce
```

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** | **per-vrf-46** | **route-policy** }
no label mode { **per-ce** | **per-vrf** | **per-vrf-46** | **route-policy** }

Use this syntax for IPv4/IPv6 AF configuration mode under vrf mode:

label mode { **per-prefix** | **per-ce** | **per-vrf** | **per-vrf-46** | **route-policy** }
no label mode { **per-prefix** | **per-ce** | **per-vrf** | **per-vrf-46** | **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-vrf-46	Specifies that the same label is used for all routes advertised from a unique VRF or global routing table.
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 4.3.1	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.
Release 24.1.1	The **per-vrf-46** keyword 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.

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

```
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
```

The example shows how to enable the per-vrf-46 label mode for global table IPv4 unicast:

```
Router# configure
Router(config)# router bgp 65550
Router(config-bgp-vrf)# address-family ipv4 unicast
Router(config-bgp-vrf)# label mode per-vrf-46
```

The example shows how to enable the per-vrf-46 label mode for global table IPv6 unicast:

```
Router# configure
Router(config)# router bgp 65550
Router(config-bgp-vrf)# address-family ipv6 unicast
Router(config-bgp-vrf)# label mode per-vrf-46
```

The example shows how to enable the per-vrf-46 label mode for global table VPNv4 unicast:

```
Router# configure
Router(config)# router bgp 65550
Router(config-bgp-vrf)# address-family vpnv4 unicast
Router(config-bgp-vrf)# vrf-all
Router(config-bgp-vrf)# label mode per-vrf-46
```

The example shows how to enable the per-vrf-46 label mode for global table VPNv6 unicast:

```
Router# configure
Router(config)# router bgp 65550
Router(config-bgp-vrf)# address-family vpnv6 unicast
Router(config-bgp-vrf)# vrf-all
Router(config-bgp-vrf)# label mode per-vrf-46
```

The example shows how to enable the per-vrf-46 label mode for VRF IPv6 unicast:

```
Router# configure
Router(config)# router bgp 65550
Router(config-bgp-vrf)# vrf INET
Router(config-bgp-vrf)# address-family ipv6 unicast
Router(config-bgp-vrf)# label mode per-vrf-46
```

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 3.7.2	This command was introduced.
Release 3.9.0	Asplain format for 4-byte Autonomous system numbers notation was supported. The **dual-as** keyword was added. The **disable** keyword was replaced with the **inheritance-disable** keyword.
Release 5.2.2	Support was added to specify the same autonomous system number for **local-as** and **remote-as** commands.

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.

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). 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.



Important If you modify the **local-as** configuration to add the **replace-as** keyword, the BGP neighbor session will reset.

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/RSP0/CPU0:router(config)# router bgp 140
RP/0/RSP0/CPU0:router(config-bgp)# neighbor 172.20.1.1
RP/0/RSP0/CPU0:router(config-bgp-nbr)# remote-as 300
RP/0/RSP0/CPU0:router(config-bgp-nbr)# local-as 30
```

Related Commands

Command	Description
neighbor-group, on page 231	Creates a neighbor group and enters neighbor group configuration mode.
session-group, on page 307	Creates a session group and enters session group configuration mode.

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 5.2.2	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	Operations
bgp	read

Examples

This example shows how to configure the neighbor to be LLGR capable for the given address family:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router bgp 100
RP/0/RSP0/CPU0:router(config-bgp)# neighbor 3.3.3.3
RP/0/RSP0/CPU0:router(config-bgp-nbr)# address-family vpnv4 unicast
RP/0/RSP0/CPU0: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 :

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router bgp 100
RP/0/RSP0/CPU0:router(config-bgp)# neighbor 3.3.3.3
RP/0/RSP0/CPU0:router(config-bgp-nbr)# address-family vpnv4 unicast
RP/0/RSP0/CPU0: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.

Related Commands

Command	Description
bgp graceful-restart, on page 91	Enables graceful restart on a BGP neighbor.
show bgp neighbors, on page 417	Displays information about BGP connections to neighbors.

lpts punt excessive-flow-trap penalty-timeout bgp

To set the penalty timeout for the bgp protocol, use the **lpts punt excessive-flow-trap penalty-timeout bgp** command in the Global Configuration mode. To restore the default penalty timeout value, use the **no** form of this command.

lpts punt excessive-flow-trap {**penalty-timeout bgp**}*timeout*
no lpts punt excessive-flow-trap {**penalty-timeout bgp**}

Syntax Description

timeout	The penalty timeout value for the bgp protocol in minutes. It is the period of time at which bgp bad flow remains to be in bad actor state. Value ranges from 1 to 1000.

Command Default

The default penalty timeout value is 15 minutes.

Command Modes

Global Configuration mode

Command History

Release	Modification
Release 6.0.1	This command was introduced.

Usage Guidelines

You must be in a user group associated with a task group that includes the proper task IDs. The command reference guides include the task IDs required for each command. If you suspect user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

If the penalty-timeout value for bgp is configured as 20, then bgp packets are dropped for 20 minutes.

Task ID

Task ID	Operations
lpts	read
basic-services	read-write

Examples

This example shows to set penalty time out for bgp bad actor:

```
RP/0/RSP0/CPU0:router(config)# lpts punt excessive-flow-trap
RP/0/RSP0/CPU0:router(config-control-plane-policer)#  penalty-timeout bgp <1-1000>
```

Related Commands

Command	Description
show running-config lpts punt excessive-flow-trap, on page 537	Displays the running configuration for the Excessive Punt Flow Trap feature.

lpts punt excessive-flow-trap routing-protocols-enable

To enable the Excessive Punt Flow Trap (EPFT) on routing protocol packets OSPF and BGP, use the **lpts punt excessive-flow-trap routing-protocol-enable** command in the Global Configuration mode. To disable, use the **no** form of this command.

lpts punt excessive-flow-trap routing-protocols-enable
no lpts punt excessive-flow-trap routing-protocols-enable

Command Default

None

Command Modes

Global Configuration mode

Command History

Release	Modification
Release 6.0.1	This command was introduced.

Usage Guidelines

You must be in a user group associated with a task group that includes the proper task IDs. The command reference guides include the task IDs required for each command. If you suspect user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

When subscriber interface or interface-based-flow is configured, you can not configure the routing-protocol-enable command. The reverse of this also holds good, that is, if the routing-protocol-enable command is configured, you cannot configure a subscriber interface or interface-based-flow.

L3 routing-protocols to be enabled on non-subscriber interfaces mac.

Note The routing-protocols-enable command can be configured only after configuring non-subscriber-interface mac.

Task ID

Task ID	Operations
lpts	read
basic-services	read-write

Examples

This example shows to enable EPFT on L3 routing protocols:

```
RP/0/RSP0/CPU0:router(config)# lpts punt excessive-flow-trap
RP/0/RSP0/CPU0:router(config-control-plane-policer)#  non-subscriber-interfaces mac
RP/0/RSP0/CPU0:router(config-control-plane-policer)#  routing-protocols-enable
```

Related Commands

Command	Description
show running-config lpts punt excessive-flow-trap, on page 537	Displays the running configuration for the Excessive Punt Flow Trap feature.

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* }
no match flow-tag {*flow-tag number* }

Syntax Description

flow-tag number	A flow-tag number. Range is from 1 to 63.

Command Default

No match criteria is specified.

Command Modes

Class-map configuration

Command History

Release	Modification
Release 5.2.2	This command was introduced.

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes 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
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** } *maximum* [**unequal-cost**] | **eibgp** *maximum* [**equal-cost**] } **route-policy** *route-policy-name* [**selective**] }

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. Range is 2 to 8
unequal-cost	(Optional) Allows iBGP multipaths to have different BGP next-hop Interior Gateway Protocol (IGP) metrics. This option is available when either the **ibgp** or **ebgp** keywords are used.
equal-cost	(Optional) Allows eiBGP multipaths to have same BGP next-hop Interior Gateway Protocol (IGP) metrics. This option is available when the **eibgp** keyword is used.
route-policy	(Optional) Specifies the route policy to select multipath.
route-policy-name	(Optional) Name of the route policy.
selective	(Optional) Allow multipaths only from marked neighbors.

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 3.7.2	This command was introduced.
Release 7.10.1	This command was modified. The **equal-cost**, **route-policy** keywords and *route-policy-name* argument were added.

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.

Use the **equal-cost** keyword to maintain consistent nexthop IGP metric criteria for multipath eligibility among iBGP path types.

See *Implementing BGP* in the *BGP Configuration Guide for ASR 9000 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-bgp)# address-family ipv4 unicast
Router(config-bgp-af)# maximum-paths ebgp 4
```

The following example shows to how to preserve consistent nexthop IGP metric criteria for multipath eligibility among similar path types, either iBGP or eBGP.

```
Router(config)# router bgp 109
Router(config-bgp)# address-family ipv4 unicast
Router(config-bgp-af)# maximum-paths eibgp 32 equal-cost route-policy EIBGP
```

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*] [**warning-only**] [**restart** *time-interval*]
no maximum-prefix *maximum* [*threshold*] [**warning-only**] [**restart** *time-interval*]

Syntax Description

maximum	Maximum number of prefixes allowed from this neighbor. Range is from1 to 4294967295.
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

When this command is not specified, the following defaults apply:

- IPv4 Unicast: 1048576
- IPv4 Labeled-unicast: 131072
- IPv6 Unicast: 524288
- IPv6 Labeled-unicast: 131072
- IPv4 Tunnel: 1048576
- IPv4 Multicast: 131072
- IPv6 Multicast: 131072
- IPv4 MVPN: 2097152
- VPNv4 Unicast: 2097152
- IPv4 MDT: 131072
- VPNv6 Unicast: 1048576

- L2VPN EVPN: 2097152
- IPv4 Flowspec: 1048576
- IPv6 Flowspec: 524288
- VPNv4 Flowspec: 2097152
- VPNv6 Flowspec: 1048576

The default threshold, when a warning message is generated, is 75 percent.



Note Starting from Cisco IOS XR Release 7.3.1, the command has been modified to remove the default maximum-prefix value. You must configure the maximum-prefix value.

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 3.7.2	This command was introduced.
Release 7.3.1	The command has been modified to remove the default maximum-prefix value.

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 4 (IPv4) unicast prefixes allowed from the neighbor 192.168.40.24 set to 1000:

```
RP/0/RSP0/CPU0:router(config-bgp)# router bgp 109
RP/0/RSP0/CPU0:router(config-bgp)# neighbor 192.168.40.24
RP/0/RSP0/CPU0:router(config-bgp-nbr)# remote-as 1
RP/0/RSP0/CPU0:router(config-bgp-nbr)#address-family ipv4 unicast
RP/0/RSP0/CPU0:router(config-bgp-nbr-af)# maximum-prefix 1000
```

Related Commands

Command	Description
af-group, on page 28	Creates an address family group for BGP neighbors and enters address family group configuration mode.
clear bgp, on page 137	Resets a BGP connection using BGP hard or soft reconfiguration.
neighbor-group, on page 231	Creates a neighbor group and enters neighbor group configuration mode.

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 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

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/RSP0/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/RSP0/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/RSP0/CPU0:router#show mpls interfaces
  Interface                  LDP      Tunnel   Enabled
  -------------------------- -------- -------- --------
  GigabitEthernet0/1/0/2     No       No       Yes
```

Related Commands

Command	Description
address-family (BGP), on page 16	Enters address family configuration mode for configuring BGP routing sessions.

multipath

Enables multiple paths for a BGP neighbor.

To disable this function, use the **no** form of this command.

multipath
no multipath

Command Default

Multipath is disabled by default.

Command Modes

Router BGP neighbor configuration

Command History

Release	Modification
Release 4.2	This command was introduced.

Usage Guidelines

To configure BGP selective multipath feature, the **multipath** option must be enabled on the required BGP neighbor. The **multipath** configuration for a neighbor works when configured with the **selective** option of the **maximum-paths** command.

Task ID

Task ID	Operations
BGP	read, write

Examples

The following example shows how to enable multiple paths for a BGP neighbor.

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router bgp 1
RP/0/RSP0/CPU0:router(config-bgp)# address-family ipv4 unicast
RP/0/RSP0/CPU0:router(config-bgp-af)# maximum-paths ibgp 4 selective
RP/0/RSP0/CPU0:router(config-bgp-af)# maximum-paths ebgp 5 selective
RP/0/RSP0/CPU0:router(config-bgp-af)# neighbor 1.1.1.1
RP/0/RSP0/CPU0:router(config-bgp-nbr)# address-family ipv4 unicast
RP/0/RSP0/CPU0:router(config-bgp-nbr-af)# multipath
RP/0/RSP0/CPU0:router(config-bgp-nbr-af)# commit
```

mvpn

To enable BGP instance to connect to PIM/PIM6, use the **mvpn** command in router configuration 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

Router configuration

Command History

Release	Modification
Release 4.2.0	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

This example shows how to configure mvpn and enable PIM/PIM6 connection:

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#router bgp 100
RP/0/RSP0/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 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

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:

```
Router(config)# router bgp 1
Router(config-bgp)# neighbor 172.168.40.24
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:

```
Router(config)# router bgp 100
Router(config-bgp)# address-family ipv6 unicast
Router(config-bgp-af)# exit
Router(config-bgp)# neighbor 3000::1
Router(config-bgp-nbr)# remote-as 2002
Router(config-bgp-nbr)# address-family ipv6 unicast
```

Related Commands

Command	Description
address-family (BGP), on page 16	Enters address family configuration mode for configuring BGP routing sessions.
remote-as (BGP), on page 277	Adds an entry to the BGP neighbor table.
use, on page 592	Inherits characteristics from a neighbor group, session group, or address family group.

neighbor-group

To create a neighbor group and enter neighbor group configuration mode, use the **neighbor-group** command in router configuration 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

Router 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

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.

```
Router(config)# router bgp 65530
Router(config-bgp)# neighbor-group group1
Router(config-bgp-nbrgrp)# remote-as 65535
Router(config-bgp-nbrgrp)# advertisement-interval 2
Router(config-bgp-nbrgrp)# address-family ipv4 unicast
Router(config-bgp-nbrgrp-af)# send-community-ebgp
Router(config-bgp-nbrgrp-af)# exit
Router(config-bgp-nbrgrp)# address-family ipv4 multicast
Router(config-bgp-nbrgrp-af)# next-hop-self
Router(config-bgp-nbrgrp-af)# exit
Router(config-bgp-nbrgrp)# exit
Router(config-bgp)#neighbor 10.0.0.1
Router(config-bgp-nbr)# use neighbor-group group1
Router(config-bgp-nbr)# exit
Router(config-bgp)# neighbor 10.0.0.2
Router(config-bgp-nbr)# use neighbor-group group1
Router(config-bgp-nbr)# exit
```

Related Commands

Command	Description
address-family (BGP), on page 16	Enters various address family configuration modes for configuring BGP routing sessions.
neighbor (BGP), on page 229	Enters neighbor configuration mode for configuring BGP routing sessions.
use, on page 592	Inherits characteristics from a neighbor group, a session group, or an address family group.

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 5.3.1	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/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router bgp 100
RP/0/RSP0/CPU0:router(config-bgp)# vrf blue neighbor 10.10.10.1
RP/0/RSP0/CPU0:router(config-bgp-vrf-nbr)# internal-vpn-client
```

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 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

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:

```
Router(config)# router bgp 109
Router(config-bgp)# address-family ipv4 unicast
Router(config-bgp-af)# network 192.168.40.0/24 backdoor
```

Related Commands

Command	Description
network (BGP), on page 236	Specifies a local network that the BGP routing process should originate and advertise to its neighbors.

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 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

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:

```
Router(config)#router bgp 120
Router(config-bgp)# address-family ipv4 unicast
Router(config-bgp-af)# network 172.20.0.0/16
```

Related Commands

Command	Description
network backdoor, on page 234	Specifies a backdoor route to a BGP border router that provides better information about the network.
redistribute (BGP), on page 271	Redistributes routes from one routing domain into another routing domain.

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

VPNv4 address family group configuration

IPv4 neighbor address family configuration

VPNv4 neighbor address family configuration

IPv4 neighbor group address family configuration

IPv6 neighbor group address family configuration

VPNv4 neighbor group address family configuration

Command History

Release	Modification
Release 3.7.2	This command was introduced.
Release 3.9.0	The **disable** keyword was replaced with the **inheritance-disable** keyword.

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.

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/RSP0/CPU0:router(config)# router bgp 140
RP/0/RSP0/CPU0:router(config-bgp)# neighbor 172.20.1.1
RP/0/RSP0/CPU0:router(config-bgp-nbr)# remote-as 1
RP/0/RSP0/CPU0:router(config-bgp-nbr)# address-family ipv4 unicast
RP/0/RSP0/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/RSP0/CPU0:router(config)# router bgp 140
RP/0/RSP0/CPU0:router(config-bgp)# af-group group1 address-family ipv4 unicast
RP/0/RSP0/CPU0:router(config-bgp-afgrp)# next-hop-self
RP/0/RSP0/CPU0:router(config-bgp-afgrp)# exit
RP/0/RSP0/CPU0:router(config-bgp)# neighbor 172.20.1.1
RP/0/RSP0/CPU0:router(config-bgp-nbr)# remote-as 1
RP/0/RSP0/CPU0:router(config-bgp-nbr)# address-family ipv4 unicast
RP/0/RSP0/CPU0:router(config-bgp-nbr-af)# use af-group group1
RP/0/RSP0/CPU0:router(config-bgp-nbr-af)# next-hop-self inheritance-disable
```

Related Commands

Command	Description
af-group, on page 28	Creates an address family group for BGP neighbors and enters address family group configuration mode.
neighbor-group, on page 231	Creates a neighbor group and enters neighbor group configuration mode.
use, on page 592	Inherits characteristics from a neighbor group, session group, or address family group.

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**]

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.

Command Default

Overwriting of the next hop is allowed.

Command Modes

VPNv4 address family group configuration

VPNv4 neighbor address family configuration

VPNv4 neighbor group address family configuration

IPv4 labeled-unicast address family configuration

IPv6 labeled-unicast address family configuration

IPv4 address family configuration

IPv6 address family configuration

Command History

Release	Modification
Release 3.7.2	This command was introduced.
Release 3.9.0	The **disable** keyword was replaced with the **inheritance-disable** keyword.
Release 4.0.0	This command was supported in the following address family configuration modes: • IPv4 labeled-unicast address family configuration • IPv6 labeled-unicast address family configuration • IPv4 unicast address family configuration • IPv6 unicast address family configuration

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/RSP0/CPU0:router(config)# router bgp 140
RP/0/RSP0/CPU0:router(config-bgp)# af-group group1 address-family ipv4 unicast
RP/0/RSP0/CPU0:router(config-bgp-afgrp)# next-hop-unchanged disable
RP/0/RSP0/CPU0:router(config-bgp-afgrp)# exit
```

Related Commands

Command	Description
next-hop-self, on page 238	Disables next-hop calculation and allows you to insert your own address in the next-hop field of BGP updates.
use, on page 592	Inherits characteristics from a neighbor group, session group, or address family group.

bgp nexthop resolution allow-default

By default, the next hop resolution in BGP does not take the default route into account. By configuring this command, the default route is used for resolving the next-hop of BGP routes. The next hop resolution is important in deciding if the next hop for a BGP route is accessible or not.

If the BGP route has an inaccessible next hop, the route does not have a best path and will not be advertised.

Warning Applying or removing this command will reset all BGP neighbors.

bgp nexthop resolution allow-default

Syntax Description

allow-default	Enable nexthops resolution using default route.

Command Default

This applies to IPv4 and IPv6. The default route is 0.0.0.0/0 for IPv4 and ::/0 for IPv6.

Command Modes

XR Config mode

Command History

Release	Modification
Release 6.2	This command was introduced.

Task ID

Task ID	Operation
bgp	read, write

The following example shows how to configure BGP with `nexthop resolution allow-default` :

```
RP/0/0/CPU0:R1(config)#router bgp 65000
RP/0/0/CPU0:R1(config-bgp)#nexthop resolution allow-default
```

"NEXTHOP configuration changed" is seen as the last reset reason with the `show bgp neighbor` command when the `nexthop resolution allow-default` command is applied or removed:

```
RP/0/0/CPU0:R1#show bgp neighbor 10.0.0.2
...
  Last reset 00:01:59, due to NEXTHOP configuration changed

RP/0/0/CPU0:R1#show bgp neighbor 2001:db8:1::2
...
  Last reset 00:02:47, due to NEXTHOP configuration changed
```

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 in VPNv4 and VRF IPv4 Unicast. Range is 0 to 128 in IPv6 Unicast.

Command Default

Nexthop resolution for minimum prefix-length is disabled.

Command Modes

VPNv4 Unicast address family

VPNv6 Unicast address family

IPv4 Unicast address family

IPv6 Unicast address family

Link-state address family

L2VPN VPLS-VPWS address family

L2VPN EVPN address family

L2VPN MSPW address family

Command History

Release	Modification
Release 7.0.2, 7.2.1	This command was supported under IPv6 Unicast address family.
Release 4.2.1	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

This example shows how to set the minimum prefix-length for nexthop resolution as *32*:

```
Router#configure
Router(config)#router bgp 100
Router(config-bgp)#address-family vpnv4 unicast
```

```
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

VRF address family configuration

Command History

Release	Modification
Release 3.7.2	This command was introduced.
Release 7.11.1	This command was modified. This command is extended to VRF address-family configuration mode.

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.

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
```

Related Commands

Command	Description
route-policy (RPL)	Defines a route policy and enters route-policy configuration mode.
show bgp nexthops, on page 438	Display statistical information about the BGP next hops.

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 3.7.2	This command was introduced.
Release 7.10.1	The default value of **critical** delay for VPN address family was changed from 0 msec to 50 msec.

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.

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

This example shows how to set the critical next-hop trigger delay to 3500 milliseconds:

```
RP/0/RSP0/CPU0:router(config)# router bgp 109
RP/0/RSP0/CPU0:router(config-bgp)# address-family ipv4 unicast
RP/0/RSP0/CPU0:router(config-bgp-af)# nexthop trigger-delay critical 3500
```

nsr (BGP)

To activate Border Gateway Protocol (BGP) nonstop routing (NSR), use the **nsr** command in BGP global configuration mode. To deactivate BGP NSR, use the **no**form of this command.

nsr
no nsr

Syntax Description

This command has no arguments or keywords.

Command Default

BGP NSR is not activated.

Command Modes

Global configuration

Command History

Release	Modification
Release 3.9.0	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.

Use the **nsr** command to enable the Border Gateway Protocol (BGP) Nonstop Routing (NSR) with Stateful Switchover (SSO). This enables all bgp peerings to maintain the BGP state to ensure continuous packet forwarding during events that could interrupt service.



Note From release 5.2.3, NSR is enabled by default.

Task ID

Task ID	Operations
bgp	read, write

Examples

The following example shows how to enable BGP NSR:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router bgp 120
RP/0/RSP0/CPU0:router(config-bgp)# nsr
```

The following example shows how to disable BGP NSR:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router bgp 120
RP/0/RSP0/CPU0:router(config-bgp)# no nsr
```

Related Commands

Command	Description
router bgp, on page 295	Configures the Border Gateway Protocol (BGP) routing process.
nsr process-failures switchover	Configures failover as a recovery action in case of process failures for active instances to switch over to a standby route processor (RP) or a standby distributed route processor (DRP) to maintain nonstop routing (NSR).
show bgp nsr, on page 447	Displays Border Gateway Protocol (BGP) nonstop routing (NSR) information.

nsr disable (BGP)

To disable Border Gateway Protocol (BGP) nonstop routing (NSR), use the **nsr disable** command in BGP global configuration mode. To re-enable BGP NSR, use the **no**form of this command.

nsr disable
no nsr disable

Syntax Description

This command has no arguments or keywords.

Command Default

BGP NSR is activated by default.

Command Modes

Global configuration

Command History

Release	Modification
Release 5.3.0	This command was introduced.

Usage Guidelines

Use the **nsr disable** command to disable Border Gateway Protocol (BGP) Nonstop Routing (NSR) with Stateful Switchover (SSO). Using the **no** form of this command enables all BGP peerings to maintain the BGP state to ensure continuous packet forwarding during events that could interrupt service.

Note In releases prior to R 5.2.3, NSR is disabled by default, and must be configured manually.

Task ID

Task ID	Operations
bgp	read, write

Examples

The following example shows how to disable BGP NSR:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router bgp 120
RP/0/RSP0/CPU0:router(config-bgp)# nsr disable
```

The following example shows how to re-enable BGP NSR:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router bgp 120
RP/0/RSP0/CPU0:router(config-bgp)# no nsr disable
```

Related Commands

Command	Description
router bgp, on page 295	Configures the Border Gateway Protocol (BGP) routing process.
nsr process-failures switchover	Configures failover as a recovery action in case of process failures for active instances to switch over to a standby route processor (RP) or a standby distributed route processor (DRP) to maintain nonstop routing (NSR).
show bgp nsr, on page 447	Displays Border Gateway Protocol (BGP) nonstop routing (NSR) information.

optimal-route-reflection

To enable the BGP optimal route reflector (ORR) feature, use the **optimal-route-reflection** command in router BGP, or neighbor configuration mode, as appropriate.

optimal-route-reflection *orr-group-name* *primary-ip-address* [*secondary-ip-address*] [*tertiary-ip-address*] *router-policy* { **addpath name** }

Syntax Description

orr-group-name	Specify the ORR group name. A maximum of 32 characters are allowed.
primary-ip-address	Specify the primary SPF root IP address. Depending on the address family configured under BGP, the SPF root IP address can be either IPv4 or IPv6. The primary SPF root IP address is the IP address of the router for which best path is calculated.
secondary-ip-address	[Optional] Specify the secondary SPF root IP address. Depending on the address family configured under BGP, the SPF root IP address can be either IPv4 or IPv6. The secondary SPF root IP address is the IP address of the nearest neighbor of the router for which best path is calculated.
tertiary-ip-address	[Optional] Specify the tertiary SPF root IP address. Depending on the address family configured under BGP, the SPF root IP address can be either IPv4 or IPv6. The secondary SPF root IP address is the IP address of the nearest neighbor of the router for which best path is calculated.
router-policy	[Optional] Enables the ORR group add path policy.
addpath name	[Optional] Specify the ORR group add path policy name.

Command Default

BGP ORR is disabled by default.

Command Modes

router BGP

neighbor configuration

Command History

Release	Modification
Release 24.2.1	The keyword *router-policy* was added.
Release 6.1.2	This command was introduced.

Usage Guidelines

Task ID

Task ID	Operation
bgp	read, write

This sample shows how to determine shortest exit point for the router with IP address 192.0.2.1, in the domain with AS number 6500, and ORR group name group1. This configuration is executed on virtual router reflector:

```
vRR# router bgp 6500
      address-family ipv4 unicast
        optimal-route-reflection group1 192.0.2.1
     commit
```

The following example shows how to configure the an ORR group add path policy:

```
router# configure
router(config)# router bgp 13
router(config-bgp)# optimal-route-reflection ipv4 orr-grp-1 1.1.1.1 route-policy
addpath-backup
router(config-bgp-router)# address-family ipv4 unicast
router(config-bgp-router-af)# additional-paths selection route-policy addpath-all
router(config-bgp-router-af)# optimal-route-reflection apply orr-grp-1
router(config-bgp-router-af)# commit
```

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 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 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

The following example shows how to configure outbound and inbound filtering criteria:

```
RP/0/RSP0/CPU0:router(config)#router bgp 6
RP/0/RSP0/CPU0:router(config-bgp)# neighbor 172.20.1.1
RP/0/RSP0/CPU0:router(config-bgp-nbr)# remote-as 1
RP/0/RSP0/CPU0:router(config-bgp-nbr)# address-family ipv4 unicast
RP/0/RSP0/CPU0:router(config-bgp-nbr-af)#orf route-policy policy_A
```

Related Commands

Command	Description
route-policy (BGP), on page 290	Applies a routing policy to updates advertised to or received from a BGP neighbor.

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 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

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/RSP0/CPU0:router(config)# router bgp 140
RP/0/RSP0/CPU0:router(config-bgp)#neighbor 172.20.1.1
RP/0/RSP0/CPU0:router(config-bgp-nbr)#remote-as 1
RP/0/RSP0/CPU0:router(config-bgp-nbr)#password clear password1
```

Related Commands

Command	Description
neighbor-group, on page 231	Creates a neighbor group and enters neighbor group configuration mode.
password-disable, on page 259	Overrides any inherited password configuration from a neighbor group or session group for BGP neighbors.
session-group, on page 307	Creates a session group and enters session group configuration mode.
timers (BGP), on page 567	Set the timers for a specific BGP neighbor.

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 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

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/RSP0/CPU0:router(config)# router bgp 140
RP/0/RSP0/CPU0:router(config-bgp)# session-group group1
RP/0/RSP0/CPU0:router(config-bgp-sngrp)# password clear password1
RP/0/RSP0/CPU0:router(config-bgp-sngrp)# exit
RP/0/RSP0/CPU0:router(config-bgp)# neighbor 172.20.1.1
RP/0/RSP0/CPU0:router(config-bgp-nbr)# remote-as 2
RP/0/RSP0/CPU0:router(config-bgp-nbr)# use session-group group1
```

```
RP/0/RSP0/CPU0:router(config-bgp-nbr)# password-disable
```

Related Commands

Command	Description
neighbor-group, on page 231	Creates a neighbor group and enters neighbor group configuration mode.
password (BGP), on page 257	Enables MD5 authentication on a TCP connection between two BGP neighbors.
session-group, on page 307	Creates a session group and enters session group configuration mode.
use, on page 592	Inherits characteristics from a neighbor group, a session group, or an address family group.

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 4.2.1	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.

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:

```
Router#configure
Router(config)#router bgp 100
Router(config-bgp)#rpki server 172.168.35.40
Router(config-bgp-rpki-server)# transport ssh port 22
Router(config-bgp-rpki-server)#username rpki-user
Router(config-bgp-rpki-server)#password rpki-ssh-pass
```

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 5.1.1	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 shows how to define permanent path for a route policy named POLICY-PERMANENT-NETWORK-IPv4:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router bgp 100
RP/0/RSP0/CPU0:router(config-bgp)# address-family ipv4 unicast
RP/0/RSP0/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:

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)

Command Default

No default behavior or values

Command Modes

Neighbor configuration

Neighbor session group configuration

Neighbor group 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

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/RSP0/CPU0:router(config)# router bgp 5
RP/0/RSP0/CPU0:router(config-bgp)# neighbor 10.1.1.1
RP/0/RSP0/CPU0:router(config-bgp-nbr)# remote-as 100
RP/0/RSP0/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 4.2.1	This command was introduced.
Release 7.4.1	This command was deprecated.

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

This example shows how to set preference value for RPKI configuration as 1:

```
Router#configure
Router(config)#router bgp 100
Router(config-bgp)#rpki server 172.168.35.40
Router(config-bgp-rpki-cache)# transport ssh port 22
Router(config-bgp-rpki-cache)#username rpki-user
Router(config-bgp-rpki-cache)#password rpki-ssh-pass
Router(config-bgp-rpki-cache)#preference 1
```

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 4.2.1	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

This example shows how to set the purge-time for RPKI cache as 30 seconds:

```
Router#configure
Router(config)#router bgp 100
Router(config-bgp)#rpki server 172.168.35.40
Router(config-bgp-rpki-server)# transport ssh port 22
Router(config-bgp-rpki-server)#username rpki-user
Router(config-bgp-rpki-server)#password rpki-ssh-pass
Router(config-bgp-rpki-server)#preference 1
Router(config-bgp-rpki-server)#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 3.7.2	This command was introduced.
Release 3.9.0	Asplain format for 4-byte Autonomous system numbers notation was supported.

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.

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/RSP0/CPU0:router(config)# router bgp 1
RP/0/RSP0/CPU0:router(config-bgp)# vrf vrf-1
RP/0/RSP0/CPU0:router(config-bgp-vrf)# rd auto
```

Related Commands

Command	Description
bgp router-id, on page 110	Configures a fixed router ID for a BGP-speaking router.
export route-target, on page 185	Configures a VRF export route-target extended community.
import route-target, on page 196	Configures a VRF import route-target extended community.

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 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

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/RSP0/CPU0:router(config)# router bgp 1
RP/0/RSP0/CPU0:router(config-bgp)# neighbor 172.20.1.1
RP/0/RSP0/CPU0:router(config-bgp-nbr)# remote-as 1
RP/0/RSP0/CPU0:router(config-bgp-nbr)# receive-buffer-size 65536 8192
```

Related Commands

Command	Description
neighbor-group, on page 231	Creates a neighbor group and enters neighbor group configuration mode.
send-buffer-size, on page 300	Sets the size of the send buffers for a BGP neighbor.
session-group, on page 307	Creates a session group and enters session group configuration mode.
socket receive-buffer-size, on page 554	Sets the size of the receive buffers for all BGP neighbors.

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*]

Enhanced Interior Gateway Routing Protocol (EIGRP)
redistribute eigrp *process-id* [**match** {**external** | **internal**}] [**metric** *metric-value*] [**route-policy** *route-policy-name*]
no redistribute eigrp *process-id* [**match** {**external** | **internal**}] [**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*

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.
eigrp	Specifies that routes are distributed from EIGRP. You must be in IPv4 unicast or multicast address family configuration mode or in VRF IPv4 address family configuration mode.
process-id	For the **eigrp** keyword, an EIGRP instance name from which routes are to be redistributed. 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 address-families. Redistribution is not allowed under VPNv4 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.

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.

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, eigrp, isis, ospf, rip,** and **static** are supported.
- IPv6 address family configuration, both unicast and multicast: **connected, eigrp, isis, ospfv3,** and **static** are supported.
- VRF IPv4 address family configuration: **connected**, **eigrp**, **ospf**, **rip**, and **static** are supported.
- VRF IPv6 address family configuration: **connected**, **eigrp**, and **static** are supported.
- BGP address family configuration: **service-layer** and **gribi** are supported.

Command History

Release	Modification
Release 3.7.2	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

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.



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.

```
Router(config)# router bgp 109
Router(config-bgp)# address-family ipv4 unicast
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
```

Related Commands

Command	Description
network (BGP), on page 236	Specifies a local network that the BGP routing process should originate and advertise to its neighbors.

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 4.2.1	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

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:

```
Router#configure
Router(config)#router bgp 100
Router(config-bgp)#rpki server 172.168.35.40
Router(config-bgp-rpki-server)# transport ssh port 22
Router(config-bgp-rpki-server)#username rpki-user
Router(config-bgp-rpki-server)#password rpki-ssh-pass
Router(config-bgp-rpki-server)#preference 1
Router(config-bgp-rpki-server)#purge-time 30
Router(config-bgp-rpki-server)#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 3.7.2	This command was introduced.
Release 3.9.0	Asplain format for 4-byte Autonomous system numbers notation was supported.

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.

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:

```
Router(config)# router bgp 1
Router(config-bgp)# session-group group2
Router(config-bgp-sngrp)# remote-as 1
Router(config-bgp-sngrp)#exit
Router(config-bgp)# neighbor 10.0.0.1
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.

```
Router(config)#router bgp 1
Router(config-bgp)# session-group group2
Router(config-bgp-sngrp)# remote-as 1
Router(config-bgp-sngrp)# exit
Router(config-bgp)# neighbor 10.0.0.1
Router(config-bgp-nbr)# use session-group group2
```

Related Commands

Command	Description
neighbor (BGP), on page 229	Enters neighbor configuration mode for configuring BGP routing sessions.

Command	Description
neighbor-group, on page 231	Creates a neighbor group and enters neighbor group configuration mode.
router bgp, on page 295	Configures the BGP routing process.
session-group, on page 307	Creates a session group and enters session group configuration mode.
use, on page 592	Inherits characteristics from a neighbor group, session group, or address family group.

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

VPNv4 neighbor address family configuration

VRF IPv4 neighbor address family configuration

VPNv4 neighbor group address family configuration

VRF IPv6 neighbor address family configuration

Command History

Release	Modification
Release 3.7.2	This command was introduced.
Release 3.9.0	The **disable** keyword was replaced with the **inheritance-disable** keyword.
Release 3.9.2	The **entire-aspath** keyword was supported.

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.

This feature is available for external BGP (eBGP) neighbors only.

When an update is passed to the external neighbor, the system drops any private autonomous system numbers. This happens irrespective of whether the autonomous system numbers are at the beginning 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/RSP0/CPU0:router(config)# router bgp 140
RP/0/RSP0/CPU0:router(config-bgp)# neighbor 172.20.1.1
RP/0/RSP0/CPU0:router(config-bgp-nbr)# remote-as 1
RP/0/RSP0/CPU0:router(config-bgp-nbr)# address-family ipv4 unicast
RP/0/RSP0/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/RSP0/CPU0:router(config)# router bgp 140
RP/0/RSP0/CPU0:router(config-bgp)# af-group group1 address-family ipv4 unicast
RP/0/RSP0/CPU0:router(config-bgp-afgrp)# remove-private-as
RP/0/RSP0/CPU0:router(config-bgp-afgrp)# exit
RP/0/RSP0/CPU0:router(config-bgp)# neighbor 172.20.1.1
RP/0/RSP0/CPU0:router(config-bgp-nbr# remote-as 1
RP/0/RSP0/CPU0:router(config-bgp-nbr)# address-family ipv4 unicast
RP/0/RSP0/CPU0:router(config-bgp-nbr-af)# use af-group group1
RP/0/RSP0/CPU0:router(config-bgp-nbr-af)# remove-private-as inheritance-disable
```

Related Commands

Command	Description
af-group, on page 28	Creates an address family group for BGP neighbors and enters address family group configuration mode.
neighbor-group, on page 231	Creates a neighbor group and enters neighbor group configuration mode.
remote-as (BGP), on page 277	Allows entries to the BGP neighbor table.

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 4.2.1	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

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:

```
Router#configure
Router(config)#router bgp 100
Router(config-bgp)#rpki server 72.168.35.40
Router(config-bgp-rpki-server)# transport ssh port 22
Router(config-bgp-rpki-server)#username rpki-user
Router(config-bgp-rpki-server)#password rpki-ssh-pass
Router(config-bgp-rpki-server)#preference 1
Router(config-bgp-rpki-server)#purge-time 30
Router(config-bgp-rpki-server)#refresh-time 30
Router(config-bgp-rpki-server)#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 3.9.0	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.

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:

```
Routers# configure
Routers#(config)# router bgp 100
Routers#(config-bgp)# address-family vpnv4 unicast
Routers#(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

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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

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/RSP0/CPU0:router(config)# router bgp 140
RP/0/RSP0/CPU0:router(config-bgp)# address-family vpnv4 unicast
RP/0/RSP0/CPU0:router(config-bgp-af)# retain route-target route-filter ft-policy-A
```

Related Commands

Command	Description
import route-target, on page 196	Configures a VRF import route-target extended community.

route-monitoring inbound post-policy

To enable the Adj-RIB-In post-policy (inbound post-policy) view for all the address families, run the **route-monitoring inbound post-policy** command in router configuration mode. Use the **no** form of the command to disable the feature.

route-monitoring inbound post-policy

Syntax Description

inbound post-policy	Inbound post-policy route-monitoring mode

Command Default

None

Command Modes

Router configuration

Command History

Release	Modification
Release 7.5.4	The syntax of this command was modified to enable the Adj-RIB-In post-policy (inbound post-policy) view for all the address families.
Release 5.2.2	This command was introduced.

Usage Guidelines

- To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.
- Adj-RIB-In Pre-policy view for all the AFIs is available by default if at least one BMP server is configured.
- Adj-RIB-In Pre-policy cannot be configured with **route-monitoring inbound post-policy** or Local-RIB.

Task ID

Task ID	Operations
bgp	read, write

The following example shows how to configure the route-monitoring inbound post-policy:

```
Router#config
Router(config)#bmp server all
Router(config-bgp-bmp)#route-monitoring inbound post-policy
Router(config-bgp-bmp-rmon)#commit
```

route-monitoring local-rib

To enable the local RIB view for all the address families, run the **route-monitoring local-rib** command in router configuration mode. Use the **no** form of the command to disable the feature.

route-monitoring local-rib

Syntax Description

local-rib	Local-RIB route-monitoring mode

Command Default

None

Command Modes

Router configuration

Command History

Release	Modification
Release 7.5.4	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

The following example shows how to configure the route-monitoring local-rib:

```
Router#config
Router(config)#bmp server all
Router(config-bgp-bmp)#route-monitoring local-rib
Router(config-bgp-bmp-rmon)#commit
```

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** }

route-policy *aggregate-route-policy-name* [**set aggregate-contributor**]

route-policy [set_comm_in | **set_comm_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.
set_comm_in	Apply the inbound policy to the neighbors after setting the BGP attributes to the aggregate-contributor.
set_comm_out	Apply the outbound policy to the neighbors after setting the BGP attributes to the aggregate-contributor.
set aggregate-contributor	Sets the specific routes as aggregate contributors.

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

VRF IPv6 neighbor address family configuration

Command History

Release	Modification
Release 3.7.2	This command was introduced.
Release 7.5.4	The **set aggregate-contributor** keyword is added.

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.

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.

The aggregate-contributor feature is applicable for the following Address Family Indicators (AFIs):

- Ipv4 unicast
- IPv6 unicast

Aggregate-contributor flag is set in the router where you configure it.

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:

```
Router(config)# router bgp 1
Router(config-bgp)# neighbor 172.20.1.1
Router(config-bgp-nbr)# remote-as 1
Router(config-bgp-nbr)# address-family ipv4 unicast
Router(config-bgp-nbr-af)# route-policy In-Ipv4 in
```

The following example shows how to set an aggregate contributor to a route policy:

```
Router# config
```

```
Router(config)#route-policy aggregate-policy1
Router(config-rpl)#set aggregate-contributor
Router(config-rpl)#end
Router(config-rpl)#commit
```

Related Commands

Command	Description
af-group, on page 28	Creates an address family group for BGP neighbors and enters address family group configuration mode.
neighbor-group, on page 231	Creates a neighbor group and enters neighbor group configuration mode.
route-policy (RPL)	Defines a route policy and enters route-policy configuration mode.

route-reflector-client

To configure the router as a Border Gateway Protocol (BGP) route reflector and configure the specified neighbor as its client, use the **route-reflector-client** command in an appropriate configuration mode. To disable configuring the neighbor as a client, use the **no** form of this command.

route-reflector-client [**inheritance-disable**]
no route-reflector-client [**inheritance-disable**]

Syntax Description

inheritance-disable	(Optional) Allows the configuration inherited from a neighbor group or address family group to be overridden.

Command Default

The neighbor is not treated as a route reflector client.

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

Command History

Release	Modification
Release 3.7.2	This command was introduced.
Release 3.9.0	The **disable** keyword was replaced with the **inheritance-disable** keyword.

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.

This command is restricted to internal BGP (iBGP) neighbors only.

Use the **route-reflector-client** command to configure the local router as the route reflector and the specified neighbor as one of its clients. All neighbors configured with this command are members of the client group, and the remaining iBGP peers are members of the nonclient group for the local route reflector.

By default, all iBGP speakers in an autonomous system must be fully meshed with each other, and neighbors do not readvertise iBGP learned routes to other iBGP neighbors.

With route reflection, all iBGP speakers need not be fully meshed. An iBGP speaker, the route reflector, passes learned iBGP routes to some number of iBGP client neighbors. Learned iBGP routes eliminate the need for each router running BGP to communicate with every other device running BGP in the autonomous system.

The local router is a route reflector as long as it has at least one route reflector client.

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 neighbor at 172.20.1.1 configured as a route reflector client for IP Version 4 (IPv4) unicast routes:

```
RP/0/RSP0/CPU0:router(config)# router bgp 140
RP/0/RSP0/CPU0:router(config-bgp)# neighbor 172.20.1.1
RP/0/RSP0/CPU0:router(config-bgp-nbr)# remote-as 140
RP/0/RSP0/CPU0:router(config-bgp-nbr)# address-family ipv4 unicast
RP/0/RSP0/CPU0:router(config-bgp-nbr-af)# route-reflector-client
```

The following example disables the route-reflector client for neighbor 172.20.1.1, preventing this feature from being automatically inherited from address family group group1:

```
RP/0/RSP0/CPU0:router(config)# router bgp 140
RP/0/RSP0/CPU0:router(config-bgp)# af-group group1 address-family ipv4 unicast
RP/0/RSP0/CPU0:router(config-bgp-afgrp)# route-reflector-client
RP/0/RSP0/CPU0:router(config-bgp-afgrp)#exit
RP/0/RSP0/CPU0:router(config-bgp)# neighbor 172.20.1.1
RP/0/RSP0/CPU0:router(config-bgp-nbr)# remote-as 140
RP/0/RSP0/CPU0:router(config-bgp-nbr)# address-family ipv4 unicast
RP/0/RSP0/CPU0:router(config-bgp-nbr-af)# use af-group group1
RP/0/RSP0/CPU0:router(config-bgp-nbr-af)# route-reflector-client inheritance-disable
```

Related Commands

Command	Description
af-group, on page 28	Creates an address family group for BGP neighbors and enters address family group configuration mode.
bgp cluster-id, on page 80	Configures the cluster ID if the BGP cluster has more than one route reflector.
neighbor-group, on page 231	Creates a neighbor group and enters neighbor group configuration mode.

router bgp

To configure the Border Gateway Protocol (BGP) routing process, use the **router bgp** command in Global Configuration mode mode. To remove all BGP configurations and terminate the BGP routing process, use the **no** form of this command.

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.

Command Default

No BGP routing process is enabled.

Command Modes

Global Configuration mode

Command History

Release	Modification
Release 3.7.2	This command was introduced.
Release 3.9.0	Asplain format for 4-byte Autonomous system number notation was supported.

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/RSP0/CPU0:router(config)# router bgp 120
```

rpki route

To statically configure an RPKI route, use the **rpki route** command in the router BGP configuration submode. The **no** form of this command removes the RPKI routes.

rpki route *ip-address-length* {**max** *max-prefix-length* | **origin** *origin-autonomous-system-number*}
no rpki route *ip-address-length* {**max** *max-prefix-length* | **origin** *origin-autonomous-system-number*}

Syntax Description

ip-address/length	Specifies the IP address of the network along with the minimum prefix length.
max *max-prefix-length*	Specifies the maximum prefix length (32 for IPv4 and 128 for IPv6).
origin *origin-autonomous-system-number*	Specifies the autonomous system number.

Command Default

RPKI route configuration is disabled.

Command Modes

Router BGP configuration

Command History

Release	Modification
Release 4.2.1	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.

In general, BGP receives the Route-Origin-Attestation (ROA) information from RPKI cache. However, the **rpki route** command is used for verification. This command can be used to configure both IPv4 and IPv6 ROAs.

This command contains all the essential attributes of an ROA record, that is, the prefix-block (IP address/length (minimum/maximum)) and the origin AS authorized to create the prefix-block.

Multiple static ROAs can be configured through this command and these entries will be included in the routers RPKI database, as if they were fetched from an RPKI cache.

Task ID

Task ID	Operation
bgp	read, write

This example shows how to configure an rpki route:

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#router bgp 100
RP/0/RSP0/CPU0:router(config-bgp)#rpki route 192.168.1.0/24 max 30 origin 65001
RP/0/RSP0/CPU0:router(config-bgp)#rpki route 172.200.0.0/16 max 24 origin 300
```

```
RP/0/RSP0/CPU0:router(config-bgp)#
```

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-name ip-address* }
no rpki server { *host-name ip-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 4.2.1	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

This example shows how to configure an rpki cache-server database and enter rpki-server configuration mode:

```
Router#configure
Router(config)#router bgp 100
Router(config-bgp)#rpki server 172.168.35.40
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 global configuration 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

global configuration

Command History

Release	Modification
Release 4.2.0	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.

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/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#selective-vrf-download disable
```

Related Commands

Command	Description
show svd role, on page 539	Displays Selective VRF Download (SVD) role information.
show svd state, on page 540	Displays Selective VRF Download (SVD) state information.

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 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

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/RSP0/CPU0:router(config)# router bgp 1
RP/0/RSP0/CPU0:router(config-bgp)# neighbor 172.20.1.1
RP/0/RSP0/CPU0:router(config-bgp-nbr)# remote-as 1
RP/0/RSP0/CPU0:router(config-bgp-nbr)# send-buffer-size 8192 8192
```

Related Commands

Command	Description
neighbor-group, on page 231	Creates a neighbor group and enters neighbor group configuration mode.
receive-buffer-size, on page 269	Sets the size of the receive buffers for a BGP neighbor.
session-group, on page 307	Creates a session group and enters session group configuration mode.
socket send-buffer-size, on page 556	Sets the size of the send buffers for all BGP neighbors.

send-community-ebgp

To specify that community or the 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 3.7.2	This command was introduced.
Release 3.9.0	The **disable** keyword was replaced with the **inheritance-disable** keyword.
Release 6.3.1	This command supports large communities.
Release 7.11.1	This command was deprecated. The command remains available only to support OpenConfig configurations and is not intended for native CLI use.

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.

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 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:

```
Router(config)#router bgp 140
Router(config-bgp)# neighbor 172.20.1.1
Router(config-bgp-nbr)# remote-as 1
Router(config-bgp-nbr)# address-family ipv4 multicast
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:

```
Router(config)#router bgp 140
Router(config-bgp)# af-group group1 address-family ipv4 multicast
Router(config-bgp-afgrp)# send-community-ebgp
Router(config-bgp-afgrp)# exit
Router(config-bgp)# neighbor 172.20.1.1
Router(config-bgp-nbr)# remote-as 1
Router(config-bgp-nbr)# address-family ipv4 multicast
Router(config-bgp-nbr-af)# use af-group group1
Router(config-bgp-nbr-af)# send-community-ebgp inheritance-disable
```

Related Commands

Command	Description
af-group, on page 28	Creates an address family group for BGP neighbors and enters address family group configuration mode.
neighbor-group, on page 231	Creates a neighbor group and enters neighbor group configuration mode.
send-extended-community-ebgp, on page 305	Specifies that extended community attributes are sent to eBGP neighbors.

send-community-gshut-ebgp

To direct the router to add the gshut community to the path having the gshut attribute or the path being sent to a connection that has graceful maintenance activated, use the **send-community-gshut-ebgp** command in the neighbor address family configuration mode. To disable the g-shut community from being announced to ebgp neighbors, use the **no** form of this command.

send-community-gshut-ebgp [**inheritance-disable**]

Syntax Description

inheritance-disable	(Optional) Prevent send-community-gshut-ebgp from being inherited from the parent.

Command Default

g-shut community attribute is not sent to eBGP 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

Command History

Release	Modification
Release 5.3.2	This command was introduced.

Usage Guidelines

Under neighbor address family configuration, use the **send-community-gshut-ebgp** command to allow the g-shut community to be sent if it is an ebgp neighbor. A path acquires the gshut attribute when it is received from a connection that has graceful maintenance activated. The sending of the gshut community if it is present because the path was received with that community or if it was added by outbound policy is governed like all other communities by the **send-community-ebgp** configuration.

Task ID

Task ID	Operations
bgp	read, write

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 3.7.2	This command was introduced.
Release 3.9.0	The **disable** keyword was replaced with the **inheritance-disable** keyword.
Release 7.11.1	This command was deprecated. The command remains available only to support OpenConfig configurations and is not intended for native CLI use.

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.

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:

```
Router(config)# router bgp 140
Router(config-bgp)# neighbor 172.20.1.1
Router(config-bgp-nbr)# remote-as 1
Router(config-bgp-nbr)# address-family ipv4 multicast
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:

```
Router(config)# router bgp 140
Router(config-bgp)# af-group group1 address-family ipv4 multicast
Router(config-bgp-afgrp)# send-extended-community-ebgp
Router(config-bgp-afgrp)# exit
Router(config-bgp)# neighbor 172.20.1.1
Router(config-bgp-nbr)# remote-as 1
Router(config-bgp-nbr)# address-family ipv4 multicast
Router(config-bgp-nbr-af)# use af-group group1
Router(config-bgp-nbr-af)# send-extended-community-ebgp inheritance-disable
```

Related Commands

Command	Description
af-group, on page 28	Creates an address family group for BGP neighbors and enters address family group configuration mode.
neighbor-group, on page 231	Creates a neighbor group and enters neighbor group configuration mode.
send-community-ebgp, on page 302	Specifies that community attributes should be sent to an eBGP neighbor.

session-group

To create a session group and enter session group configuration mode, use the **session-group** command in router configuration 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

Router 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

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.

```
Router(config)# router bgp 1
Router(config-bgp)# session-group group1
Router(config-bgp-sngrp)# remote-as 1
Router(config-bgp-sngrp)# advertisement-interval 2
Router(config-bgp-sngrp)# password password1
Router(config-bgp-sngrp)# exit
Router(config-bgp)# neighbor 10.0.0.1
Router(config-bgp-nbr)# use session-group group1
Router(config-bgp-nbr)# exit
Router(config-bgp)# neighbor 10.0.0.2
Router(config-bgp-nbr)# use session-group group1
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 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

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/RSP0/CPU0:router(config)# router bgp 1
RP/0/RSP0/CPU0:router(config-bgp)# neighbor 45.67.89.01
RP/0/RSP0/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 5.2.2	This command was introduced.

Usage Guidelines

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

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/RSP0/CPU0:router(config)# route-policy policy_1
RP/0/RSP0/CPU0:router(config-rpl)# set flow-tag 12
RP/0/RSP0/CPU0:router(config-rpl)# end-policy
```

show bgp

To display entries in the Border Gateway Protocol (BGP) routing table, use the **show bgp** command in EXEC mode.

show bgp [**ipv4** { **unicast** | **multicast** | **labeled-unicast** | **all** | **tunnel** | | **flowspec** } | **ipv6** { **unicast** | **multicast** | **all** | **labeled-unicast** | | **flowspec** } | **all** {**unicast** | **multicast** | **all** | **labeled-unicast** | **mdt** | **tunnel**} | **vpnv4** { **unicast**} [**rd** *rd-address*] | **vrf** {*vrf-name* | **all**} [**ipv4** {**unicast** | **labeled-unicast**} | **ipv6** **{unicast | flowspec}**] | { **flowspec** | **unicast** } | [**instance**] | [**instances**] | **flowspec**] [*ip-address* [{*mask* | */prefix-length*} [**longer-prefixes** | **unknown-attributes** | **bestpath-compare**]]] [**standby**] [**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.
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.
flowspec	Displays flowspec configuration 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

EXEC mode

Command History

Release	Modification
Release 3.7.2	This command was introduced.
Release 3.9.0	Asplain format for 4-byte Autonomous system numbers notation was supported. The input parameters and output were modified to display 4-byte autonomous system numbers and extended communities in either asplain or asdot notations.
Release 5.3.2	The command output was modified to include graceful maintenance feature information.
Release 6.3.1	The command output was modified to include large community in the list of attributes.
Release 7.9.1	The command output was modified to include the aggregate-contributor feature examples.

Usage Guidelines

set default-afi



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 *System Management Command Reference for Cisco ASR 9000 Series Routers* 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.

Task ID

Task ID	Operations
bgp	read

Examples

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 EXEC mode:

```
RP/0/RSP0/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 3: 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.
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.

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.
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/RSP0/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

 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
```

This table describes the significant fields shown in the display.

Table 4: show bgp prefix length Field Descriptions

Field	Description
BGP router table entry	Network that is being displayed.
Versions	List of the network versions in each BGP process.
Process	Name of the BGP process.
bRIB/RIB	Version of the network for sending to the RIB. You can compare this version with the bRIB/RIB version for the process (at the top of **show bgp summary**) to verify whether the network has been sent to the RIB.
SendTblVer	Version of the network for advertising to neighbors. This can be compared with the neighbor version to determine whether the network has been advertised to a particular neighbor.
Paths	List of paths for the network (that is, routes to reach the network). The number of paths and the index of the best path are given.
not advertised to any peer	Best path was received with a NO_ADVERTISE community and is not advertised to any neighbor.
not advertised to EBGP peer	Best path was received with a NO_EXPORT community and is not advertised to any eBGP neighbor.
not advertised outside local AS	Best path was received with a LOCAL_AS community and is not advertised to peers outside the local AS.

Field	Description
Advertisements of this net are suppressed by an aggregate	Network is a more-specific prefix of a configured aggregate and has been suppressed. It is not advertised to any neighbors unless they have an unsuppress-map configured.
Advertised to update-groups	List of update-groups to which the net has been advertised. Update-groups that have only one peer are not listed here.
Advertised to peers	List of neighbors to which the net has been advertised to. Neighbors that are in one of the update-groups listed above are not listed separately. Only neighbors that are in unique update-groups are listed.
Received by speaker 0	BGP process where the path originated. This is always “speaker 0” for standalone mode. It will be the speaker-id when BGP is in distributed mode.
AS Path	Autonomous system (AS) path that was received for the path. If the AS path is empty, then “Local” is displayed. This is the case for paths that are locally generated on this router or on a neighboring router within the same AS.
aggregated by	If the path is an aggregate, the router-id of the router that performed the aggregation.
suppressed due to dampening	Path has been suppressed due to the configured path dampening.
history entry	Path is withdrawn, but a copy is kept to store the dampening information.
Received from a RR-client	Path was received from a route reflector client.
received-only	If soft reconfiguration inbound is configured, the path was received but dropped by inbound policy, or was accepted and modified. In either event, the received-only value is a copy of the original, unmodified path.
received & used	If soft reconfiguration inbound is configured, the path was received and accepted by inbound policy, but not modified.
stale	Neighbor from which the path was received is down, and the path is kept and marked as stale to support graceful restart.
<nexthop> from <neighbor> (<router-id>)	Next hop for the path. If the next hop is known by a mechanism outside BGP (for example, for redistributed paths), then 0.0.0.0 is displayed. After the next hop, the neighbor from whom the path was received is displayed, along with the neighbor’s router-id. If the path was locally generated (for example, an aggregate or redistributed path), then 0.0.0.0 is displayed for the neighbor address.
Origin	IGP: the path originated from an IGP. EGP: the path originated from an EGP. incomplete: the origin of the path is unknown.
metric	MED value of the path.

Field	Description
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.
weight	Locally assigned weight (if not 0) of the path. Weight is used in choosing the preferred path to a route. It is not advertised to any neighbor.
valid	Path is valid and can be considered in the best-path calculation.
redistributed	Path is redistributed through a **redistribute** command.
aggregated	Path is a locally generated aggregate created due to an **aggregate-address** command.
local	Path is a local network source due to a **network** command.
internal	Path was received from an iBGP neighbor.
external	Path was received from an eBGP neighbor.
atomic-aggregate	Path was received with the atomic-aggregate flag set. Some path information has been removed through aggregation.
best	Path is the best path for the network and is used for routing and advertised to peers.
multipath	Path is a multipath and is installed into the RIB along with the best path.
Community	List of communities attached to the path.
Extended community	List of extended communities attached to the path.
Originator	Originator of the path within the AS Cluster list if the path is reflected.
AS Cluster list	List of RR clusters the path has passed through if the path is reflected
Dampinfo	Penalty and reuse information if the path is dampened.
penalty	Current penalty for the path.
flapped	Number of times the path has flapped and the time since the first flap.
reuse in	Time until the path is re-used (undampened).
half life	Configured half-life for the path.
suppress value	Penalty at which the path is suppressed.
reuse value	Penalty at which the path is re-used.
Maximum suppress time	Maximum length of time for which the path can be suppressed.

The following is sample output from the **show bgp** command with the *ip-address/prefix-length* **detail** options:

```
RP/0/RSP0/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
```

This is sample output to explain 'import suspect' state and 'import-suspect' field in **show bgp** command output:

```
RP/0/RSP0/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'.

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 a sample output of **show bgp** command that displays large-communities in the list of attributes:

```
RP/0/RSP0/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 is a sample output of **show bgp** command that displays the net as an aggregate-contributor after applying the aggregate routing policy:

```
Router#show bgp 250.2.2.1/32
Thu Dec  1 10:12:15.374 EST
BGP routing table entry for 250.2.2.1/32
Versions:
  Process           bRIB/RIB  SendTblVer
  Speaker                247         247
Last Modified: Dec  1 09:00:20.000 for 01:11:55
Paths: (1 available, best #1)
Net is an aggregate-contributor
  Advertised IPv4 Unicast paths to update-groups (with more than one peer):
    0.2
  Advertised IPv4 Unicast paths to peers (in unique update groups):
    192.168.0.5     10.10.10.1
  Path #1: Received by speaker 0
```

```
  Advertised IPv4 Unicast paths to update-groups (with more than one peer):
    0.2
  Advertised IPv4 Unicast paths to peers (in unique update groups):
    192.168.0.5     10.10.10.1
  105
    13.0.1.1 from 13.0.1.1 (13.0.1.1)
      Origin IGP, localpref 100, valid, external, best, group-best
      Received Path ID 0, Local Path ID 1, version 247
      Community: 20:20
      Origin-AS validity: (disabled)
```

This is a sample output of **show bgp** command that displays all aggregate-contributors of the BGP address 250.2.2.0/24:

```
Router1#show bgp 250.2.2.0/24 aggregate-contributors
Thu Dec  1 10:15:20.393 EST
BGP router identifier 192.168.0.2, local AS number 100
BGP generic scan interval 60 secs
Non-stop routing is enabled
BGP table state: Active
Table ID: 0xe0000000   RD version: 247
BGP main routing table version 247
BGP NSR Initial initsync version 22 (Reached)
BGP NSR/ISSU Sync-Group versions 247/0
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
*> 250.2.2.1/32       13.0.1.1                               0 105 i
*> 250.2.2.2/32       13.0.2.1                               0 105 i
*> 250.2.2.3/32       13.0.3.1                               0 101 i
*> 250.2.2.4/32       13.0.4.1                               0 101 i
*> 250.2.2.5/32       13.0.5.1                               0 102 i
*> 250.2.2.6/32       13.0.6.1                               0 102 i
*> 250.2.2.7/32       13.0.7.1                               0 103 i
*> 250.2.2.8/32       13.0.8.1                               0 103 i
*> 250.2.2.9/32       13.0.9.1                               0 104 i
*> 250.2.2.10/32      13.0.10.1                              0 104 i
*> 250.2.2.11/32      0.0.0.0                  0         32768 ?
*> 250.2.2.12/32      0.0.0.0                  0         32768 ?
*> 250.2.2.13/32      0.0.0.0                  0         32768 ?

Processed 13 prefixes, 13 paths
```

After setting an aggregate-contributor to a specific aggregate-address, you can then set the BGP attributes (for example, cost community, next-hop, BGP multiple exit discriminator) to the aggregate-contributor, and then apply the inbound policy and the outbound policy to the neighbors. This is a sample output of **show bgp** command that displays these details:

For inbound policy:

```
Router1#show bgp 250.2.2.1/32
Thu Dec  1 10:12:15.374 EST
BGP routing table entry for 250.2.2.1/32
Versions:
  Process           bRIB/RIB  SendTblVer
  Speaker               247          247
Last Modified: Dec  1 09:00:20.000 for 01:11:55
Paths: (1 available, best #1)
```

```
Net is an aggregate-contributor
  Advertised IPv4 Unicast paths to update-groups (with more than one peer):
    0.2
  Advertised IPv4 Unicast paths to peers (in unique update groups):
    192.168.0.5     10.10.10.1
  Path #1: Received by speaker 0
  Advertised IPv4 Unicast paths to update-groups (with more than one peer):
    0.2
  Advertised IPv4 Unicast paths to peers (in unique update groups):
    192.168.0.5     10.10.10.1
  105
    13.0.1.1 from 13.0.1.1 (13.0.1.1)
      Origin IGP, localpref 100, valid, external, best, group-best
      Received Path ID 0, Local Path ID 1, version 247
      Community: 20:20
      Origin-AS validity: (disabled)
```

For outbound policy:

```
Router2# show bgp ipv4 u 250.2.2.4/32
Thu Dec  8 15:08:40.672 EST
BGP routing table entry for 250.2.2.4/32
Versions:
  Process           bRIB/RIB  SendTblVer
  Speaker                174          174
Last Modified: Dec  8 15:06:53.000 for 00:01:47
Paths: (1 available, best #1)
  Advertised IPv4 Unicast paths to update-groups (with more than one peer):
    0.1
  Advertised IPv4 Unicast paths to peers (in unique update groups):
    12.12.12.4
  Path #1: Received by speaker 0
  Advertised IPv4 Unicast paths to update-groups (with more than one peer):
    0.1
  Advertised IPv4 Unicast paths to peers (in unique update groups):
    12.12.12.4
  101, (Received from a RR-client)
    13.0.4.1 (metric 2) from 192.168.0.2 (192.168.0.2)
      OC-RIB Attribute-Index 0
      Origin IGP, localpref 100, valid, internal, best, group-best
      Received Path ID 1, Local Path ID 1, version 174
      Extended community: RT:200:200
```

This example shows a sample output of **show bgp ipv4 flowspec** and **show bgp ipv6 flowspec** commands.

```
Router# show bgp ipv4 flowspec
Wed Apr 23 18:27:57.484 PDT
BGP router identifier 192.0.2.2, local AS number 64540
BGP generic scan interval 60 secs
Non-stop routing is enabled
BGP table state: Active
Table ID: 0x0   RD version: 972
BGP table nexthop route policy:
BGP main routing table version 972
BGP NSR Initial initsync version 2 (Reached)
BGP NSR/ISSU Sync-Group versions 0/0
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
*>iDest:209.165.201.5/27/48
```

```
                     10.1.3.2                       100      0 i
* i                   10.1.3.2                       100      0 i
-snip-
*>iSource:209.165.201.2/27/40
                     10.1.2.1                       100      0 i
* i                   10.1.2.1                       100      0 i

Processed 240 prefixes, 480 paths
```

```
Router# show bgp ipv6 flowspec
Wed Apr 23 18:28:13.518 PDT
BGP router identifier 192.0.2.2, local AS number 64540
BGP generic scan interval 60 secs
Non-stop routing is enabled
BGP table state: Active
Table ID: 0x0   RD version: 242
BGP table nexthop route policy:
BGP main routing table version 242
BGP NSR Initial initsync version 2 (Reached)
BGP NSR/ISSU Sync-Group versions 0/0
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
*>iDest:2001:DB8::3/0-128/152
                     2001:DB8::2                       100      0 i
* i                   2001:DB8::2                       100      0 i
-snip-
*>iSource:2001:DB8::1/0-64/88
                     2001:DB8::3                       100      0 i
* i                   2001:DB8::3                       100      0 i

Processed 60 prefixes, 120 paths
```

This table describes the significant fields shown in the display.

Table 5: show bgp ipv4/ipv6 flowspec field descriptions

Field	Description
Network	Indicates the IPv4 or IPv6 FlowSpec prefix or route being advertised. **IPv4:** In the example **209.165.201.5/27/48**, • **209.165.201.5/27/48** indicates the IPv4 prefix. • **/27** indicates the netmask, meaning the first 27 bits of the IP address define a single, specific host address. This netmask helps specify the exact range of IP addresses that the flowspec rule applies to. • **/48** indicates the Network Layer Reachability Information (NLRI) length, representing the total length of the FlowSpec NLRI including match fields. The format is: `<match fields>:<prefix>/<netmask>/<NLRI length>` **IPv6:** In the example **2001:DB8::1/0-64/88**, • **2001:DB8::** indicates the IPv6 prefix. • **0-64** indicates the offset and prefix length in the FlowSpec match fields. • **/88** indicates the NLRI length, representing the total length of the FlowSpec NLRI including match fields. The format is: `<match fields>:<prefix>/<offset>-<prefix length>/<NLRI length>`

Related Commands

Command	Description
aggregate-address, on page 30	Creates an aggregate entry in a BGP routing table.
bgp default local-preference, on page 88	Changes the default local preference value.
network (BGP), on page 236	Specifies a local network that the BGP routing process should originate and advertise to its neighbors.
route-policy (BGP), on page 290	Applies a routing policy to updates advertised to or received from a BGP neighbor.
set default-afi	Sets the default Address Family Identifier (AFI) for the current session.
set default-safi	Sets the default subaddress Family Identifier (SAFI) for the current session.
show bgp cidr-only, on page 370	Displays routes with nonnatural netmasks.

Command	Description
show bgp community, on page 374	Displays routes belonging to the specified communities.
show bgp inconsistent-as, on page 391	Displays networks with inconsistent origin autonomous system.
show bgp regexp, on page 482	Displays routes matching an AS path regular expression.
show bgp route-policy, on page 486	Displays networks that match a route policy.
show bgp summary, on page 496	Displays the status of all BGP connections.
show bgp truncated-communities, on page 507	Displays networks with community lists truncated by policy.

show bgp advertised

To display advertisements for neighbors or a single neighbor, use the **show bgp advertised**command in EXECEXEC mode mode.

show bgp [**ipv4** { **all** | **labeled-unicast** | **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

EXECEXEC mode

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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

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 *System Management Command Reference for Cisco ASR 9000 Series Routers* 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 EXECEXEC mode mode:

```
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 6: 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:

```
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 7: 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.

Related Commands

Command	Description
set default-afi	Sets the default Address Family Identifier (AFI) for the current session.
set default-safi	Sets the default subaddress Family Identifier (SAFI) for the current session.
route-policy (BGP), on page 290	Applies a route policy to incoming and outgoing routes.
rd, on page 267	Filters routes using a prefix list.
show bgp policy, on page 454	Displays information about BGP advertisements under a proposed policy.
sent-advertisements	Previews advertisements to peers, including details of advertised attributes.

show bgp as-neighbors

To display BGP neighbor or peer relationships that are grouped by AS, use the **show bgp as-neighbors** command.

show bgp as-neighbors [*as-number* [**standby**]]

Syntax Description

as-neighbors	Displays neighbors grouped by an ASN.
as-number	(Optional) Filters and displays the neighbors for the specified AS numbers. The range is 1-4294967295.
standby	Displays the BGP neighbor information that is maintained on the standby route processor.

Command Default

No default behavior.

Command Modes

EXEC mode and XR EXEC mode.

Command History

Release	Modification
Release 25.1.1	This command was introduced.

Usage Guidelines

None.

Task ID

Task ID	Operations
bgp	read

Examples

The following is the sample output of the **show bgp as-neighbors** command:

```
Router#show bgp as-neighbors
Wed Nov 20 21:10:58.133 UTC
AS: 4294967291, Neighbors: 64, Established: 0
Last updated: Nov 20 00:25:35.285
Neighbor                                        State                 Last state change
31.0.2.2                                        Idle                  Nov 20 19:33:34.119
…
31.0.65.2                                       Idle                  Nov 20 19:33:34.025


AS: 4294967292, Neighbors: 64, Established: 0
Last updated: Nov 20 00:25:35.285
Neighbor                                        State                 Last state change
32.0.2.3                                        Active                Nov 20 19:33:39.613
…
32.0.65.3                                       Active                Nov 20 19:33:39.540
```

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 EXECEXEC mode 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

EXECEXEC mode

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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

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 EXECEXEC mode 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.

```
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:

```
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 8: show bgp af-group Field Descriptions, on page 337 describes the significant fields shown in the display.

This table describes the significant fields shown in the display.

Table 8: 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.

Related Commands

Command	Description
af-group, on page 28	Configures a BGP address family group.
show bgp neighbors, on page 417	Displays information about BGP neighbors, including configuration inherited from neighbor groups, session groups, and address family groups.
show bgp neighbor-group, on page 413	Displays information about configuration for neighbor groups.
use, on page 592 **af-group**	Configures an af-group to inherit the configuration of a specified af-group.

show bgp aggregate-contributors

To display whether an aggregate-address has aggregate-contributors, run the **show bgp aggregate-contributors** command in EXEC mode.

show bgp *address/mask-length* **aggregate-contributors**

Syntax Description

aggregate-contributors	Displays contributors of an aggregate-address

Command Default

None

Command Modes

EXEC mode

Command History

Release	Modification
Release 7.5.4	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.
- The command is applicable for the following Address Family Indicators (AFIs):
  - IPv4 unicast
  - IPv6 unicast

Task ID

Task ID	Operation
bgp	read

The following output displays all aggregate-contributors of the BGP address 250.2.2.0/24:

```
Router1#show bgp 250.2.2.0/24 aggregate-contributors
Thu Dec  1 10:15:20.393 EST
BGP router identifier 192.168.0.2, local AS number 100
BGP generic scan interval 60 secs
Non-stop routing is enabled
BGP table state: Active
Table ID: 0xe0000000   RD version: 247
BGP main routing table version 247
BGP NSR Initial initsync version 22 (Reached)
BGP NSR/ISSU Sync-Group versions 247/0
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
*> 250.2.2.1/32       13.0.1.1                               0 105 i
*> 250.2.2.2/32       13.0.2.1                               0 105 i
```

```
*> 250.2.2.3/32       13.0.3.1                               0 101 i
*> 250.2.2.4/32       13.0.4.1                               0 101 i
*> 250.2.2.5/32       13.0.5.1                               0 102 i
*> 250.2.2.6/32       13.0.6.1                               0 102 i
*> 250.2.2.7/32       13.0.7.1                               0 103 i
*> 250.2.2.8/32       13.0.8.1                               0 103 i
*> 250.2.2.9/32       13.0.9.1                               0 104 i
*> 250.2.2.10/32      13.0.10.1                              0 104 i
*> 250.2.2.11/32      0.0.0.0                  0         32768 ?
*> 250.2.2.12/32      0.0.0.0                  0         32768 ?
*> 250.2.2.13/32      0.0.0.0                  0         32768 ?

Processed 13 prefixes, 13 paths
```

show bgp attribute-key

To display all existing attribute keys, use the **show bgp attribute-key**command in EXEC mode.

show bgp {**ipv4** | **ipv6** | **all** | **vpnv4 unicast** | **vrf**} **attribute-key**

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

EXEC

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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

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 *System Management Command Reference for Cisco ASR 9000 Series Routers* 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 EXEC mode:

```
RP/0/RSP0/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 9: 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.

Related Commands

Command	Description
set default-afi	Sets the default Address Family Identifier (AFI) for the current session.
set default-safi	Sets the default Subaddress Family Identifier (SAFI) for the current session.

show bgp bmp

To display Border Gateway Protocol (BGP) Monitoring Protocol (BMP) information, use the **show bgp bmp** command in EXEC mode.

show bgp bmp { **server** *server-id* **[detail]** | **summary** }

show bgp bmp server *<server-id>* **[[vrf** *<vrf>* **] neighbor** *<neighbor-ip>* **] [detail | wide]**

show bgp bmp [mode {inbound post-policy | local-rib | all}] *[<afi> <safi>]* **table [rd** *<rd>* **| vrf** *<vrf-name>* **] [neighbor** *<address>* **| version begin** *<begin-version>* **end** *<end-version>* **| prefix** *<prefix-length>* **| summary]**

show bgp bmp [mode {inbound post-policy | local-rib | all}] vrf *<vrf-name> [<afi> <safi>]* **table [neighbor** *<address>* **| version begin** *<begin-version>* **end** *<end-version>* **| prefix]** *<prefix-length>*

show bgp bmp mode local-rib [vrf *<vrf-name>* **] emulated-peers** *<emulated-peer-address>* **[detail | performance-statistics | brief]**

show bgp bmp mode local-rib [vrf *<vrf-name>* **] internal update-group** *<update-group-index>* **[summary] | emulated-peer** *<emulated-peer-address>* **] [performance-statistics]**

show bgp bmp mode local-rib internal update out [emulated-peer *emulated-peer-address* **]**

show bgp bmp [mode {inbound post-policy | local-rib | all}] *[<afi> <safi>]* **convergence**

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.
wide	Wide tabular formatting of the BMP neighbor details.
table	Show prefixes and paths from the Inbound Post-Policy and Local RIB tables.
ipv4	IP Version 4 address prefixes
ipv6	IP Version 6 address prefixes
vpnv4	VPN version 4 address family
vpnv6	VPN version 6 address family
unicast	Unicast sub-address family
labeled-unicast	Labeled Unicast sub-address family
all	• When used as address family, the command is executed for all address families. • When used as sub-address family, the command is executed for all sub-address families.

neighbor	Output of the command is filtered for a specific neighbor.
mode	Execute the command for the specified BMP monitoring mode.
rd	Output of the command is filtered for a specific route distinguisher.
vrf	Output of the command is filtered for a specific VRF.
summary	Show only the summary information
begin	Begin version of a version range that is used for filtering the prefixes within that range.
end	End version of a version range that is used for filtering the prefixes within that range.
emulated-peers	BMP Local-RIB emulated peers
brief	Display the output in brief tabular format.
detail	Display detailed information
performance-statistics	Display performance statistics

Command Default

No default behavior or values

Command Modes

EXEC

Command History

Release	Modification
Release 24.1.1	The option to specify AFIs and SAFIs after the **table** keyword is completely removed from Cisco IOS XR Release 24.1.1.
Release 7.5.4	The syntax of the **show bgp bmp table vpnv4 unicast** command was modified. Specifying AFI and SAFI after the **table** keyword option is not supported from Release 7.5.4. So, it is recommended to specify AFIs and SAFIs before the **table** keyword.
Release 7.5.4	This command was modified with Inbound Post-Policy and Local RIB view.
Release 5.2.2	This command was introduced.

Usage Guidelines

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

Inbound post-policy and Local-RIB tables

- If AFI is not specified while running the commands, the default IPv4 Unicast AFI is considered.
- If mode is not specified, the default mode inbound post-policy is considered.

- Specifying AFIs and SAFIs after the **table** keyword option is not supported from Release 7.5.4. So, you can specify the same before the **table** keyword as follows:

  show bgp bmp [mode <bmp-mode>] [<afi> <safi>] table

- Starting from Release 7.5.4, if AFIs or SAFIs are specified before and after the **table** keyword, the AFIs or SAFIs after the **table** keyword are considered.
- The following filters are only available for **vpnv4 unicast** and **vpnv6 unicast** AFIs:
  - **rd**
  - **vrf**
- The *begin-version* of the **bmp table version** can be smaller than or greater than or equal to *end-version*.

Emulated-Peers

- The keyword **brief** can be used only when **vrf** is not specified.
- The keyword **emulated-peers** is only applicable for Local-RIB.

BGP BMP Convergence

Use the **show bgp bmp 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:

- Table copy is not pending (applicable only in case of Inbound Post-Policy)
- Update generation is not pending
- All BMP servers that are not administratively shutdown are *Up* state
- All BMP servers have caught up to the table version
- Write queues of all BMP servers are empty

If all the above conditions are true, the address family is considered converged.

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/RSP0/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/RSP0/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
```

The **show bgp bmp summary** command displays BMP server summary information such as server address and port number, server state, time for which server is up and the number of neighbors the server is monitoring.

```
Router#show bgp bmp summary

ID   Host                 Port     State   Time        NBRs

 1   12.1.2.1             16001    ESTAB   01:35:05    14

 2   12.2.2.1             16002    ESTAB   01:30:47    11
```

The **show bgp bmp server <ID>** command displays BMP server details, BMP update modes exported to the BMP server, configured (or default) values specific to the update modes, update generation details, TCP connection status, BMP server specific message statistics, and statistics of the neighbors monitored by the BMP server.

```
Router#show bgp bmp server 1
BMP server 1
Host 12.1.2.1 Port 16001
Connected for 01:19:20
Last Disconnect event received : 00:00:00
Precedence:  internet
BGP neighbors: 14
VRF: - (0x60000000)
Update Source: - (-)
Update Source Vrf ID: 0x0
Update Mode                      : In-Post-Policy, Local-RIB
  In-Post-Policy
   Advertisement interval        : 15 secs
```

```
   Scanner interval               : 60 secs
  Local-RIB
   Advertisement interval         : 15 secs
   Scanner interval
     Global                       : 60 secs
     IPv4 Unicast                 : 60 secs
     VPNv4 Unicast                : 60 secs
     IPv6 Unicast                 : 60 secs
     VPNv6 Unicast                : 60 secs
Flapping Delay                    : 300 secs
Initial Delay                     : 0 secs
Initial Refresh Delay             : 1 secs
Initial Refresh Spread            : 1 secs
Stats Reporting Period            : 0 secs
Queue Route Mon Msg buffer limit  : 133693 KB (Current Server Up Count: 2)
Queue Route Mon Msg buffer usage  : 0 B
Queue write pulse sent            : Feb  6 12:32:13.676, not set (all)
Queue write pulse received        : Feb  6 12:32:13.676
Update Generation in Progress     : No
Reset Walk in Progress            : No
IPv4 Unicast
  Version                     : 11
  Init EOR Version            : 11
  Init EOR Pending count      : 0
  Update Generation
    Last Run                   : Feb  6 12:32:04.615, Count 1
    Walk Currently Stalled     : No, Last Stalled : not set, Count 0
VPNv4 Unicast
  Version                     : 14
  Init EOR Version            : 14
  Init EOR Pending count      : 0
  Update Generation
    Last Run                   : Feb  6 12:32:04.615, Count 1
    Walk Currently Stalled     : No, Last Stalled : not set, Count 0
IPv6 Unicast
  Version                     : 12
  Init EOR Version            : 12
  Init EOR Pending count      : 0
  Update Generation
    Last Run                   : Feb  6 12:32:04.615, Count 1
    Walk Currently Stalled     : No, Last Stalled : not set, Count 0
VPNv6 Unicast
  Version                     : 16
  Init EOR Version            : 16
  Init EOR Pending count      : 0
  Update Generation
    Last Run                   : Feb  6 12:32:04.615, Count 1
    Walk Currently Stalled     : No, Last Stalled : not set, Count 0

TCP:
  Last message sent: Feb  6 12:32:04.615, Status: No Pending Data
  Last write pulse received: Feb  6 12:32:04.863, Waiting: FALSE

Message Stats:
Total msgs dropped   : 0
Total msgs pending   : 0, Max: 19 at Feb  6 12:32:04.615
Total messages sent  : 100
Total bytes sent     : 13148, Time spent: 0.000 secs
  INITIATION         : 1
  TERMINATION        : 0
  STATS-REPORT       : 0
  PER-PEER messages  : 99
  ROUTE-MON messages : 88
   EOR messages      : 30
```

```
   Update messages   : 58 (Prefixes: 41, Err: 0)
Update gen time spent: 0.000 secs

Neighbor 12.100.0.1 (vrf foo)
  Messages pending : 0
  Messages dropped : 0
  Messages sent    : 5
    PEER-UP        : 1
    PEER-DOWN      : 0
    ROUTE-MON      : 4
      EOR          : 2
      Update       : 2 (Prefixes: 4, Err: 0)

Neighbor 12:100::1 (vrf foo)
  Messages pending : 0
  Messages dropped : 0
  Messages sent    : 3
    PEER-UP        : 1
    PEER-DOWN      : 0
    ROUTE-MON      : 2
      EOR          : 1
      Update       : 1 (Prefixes: 2, Err: 0)

Neighbor 12.101.0.1 (vrf bar)
  Messages pending : 0
  Messages dropped : 0
  Messages sent    : 5
    PEER-UP        : 1
    PEER-DOWN      : 0
    ROUTE-MON      : 4
      EOR          : 2
      Update       : 2 (Prefixes: 4, Err: 0)

Neighbor 12.0.0.1 (vrf default)
  Messages pending : 0
  Messages dropped : 0
  Messages sent    : 7
    PEER-UP        : 1
    PEER-DOWN      : 0
    ROUTE-MON      : 6
      EOR          : 4
      Update       : 2 (Prefixes: 4, Err: 0)
```

The **show bgp bmp server <ID>** command with the **wide** command displays the same information as **show bgp bmp server <ID>** except that it displays the neighbor information in a wide tabular format.

```
Router#show bgp bmp server 1 wide
BMP server 1
Host 12.1.2.1 Port 16001
Connected for 01:00:49
Last Disconnect event received : 00:00:00
Precedence:  internet
BGP neighbors: 14
VRF: - (0x60000000)
Update Source: - (-)
Update Source Vrf ID: 0x0
..
<truncated> ..

Message Stats:
Total msgs dropped   : 0
Total msgs pending   : 0, Max: 19 at Feb  6 12:32:04.615
```

```
Total messages sent  : 100
Total bytes sent     : 13148, Time spent: 0.000 secs
  INITIATION         : 1
  TERMINATION        : 0
  STATS-REPORT       : 0
  PER-PEER messages  : 99
  ROUTE-MON messages : 88
   EOR messages      : 30
   Update messages   : 58 (Prefixes: 41, Err: 0)

Update gen time spent: 0.000 secs
 Legend:
  MsgPend   : Messages yet to be sent
  MsgDrop   : Messages dropped
  MsgSent   : Messages sent
  PeerUp    : Peer-Up messages sent
  PeerDown  : Peer-Down messages sent
  RouteMon  : Router-Monitoring messages sent
  EOR       : EOR messages sent
  Update    : Update messages sent
  UpdPaths  : Paths advertised in Update messages
  UpdErrs   : Paths not advertised in Update messages
  Withdraw  : Withdraw messages sent
  WdrwPaths : Paths advertised in Withdraw messages
  WdrwErrs  : Paths not advertised in Withdraw messages

Neighbor                VRF                     MsgPend    MsgDrop    MsgSent     PeerUp
  PeerDown   RouteMon        EOR     Update   UpdPaths    UpdErrs   Withdraw  WdrwPaths
WdrwErrs

12.100.0.1              foo                       0          0          5          1
         0          4          2          2          4          0          0          0
       0
12:100::1               foo                       0          0          3          1
         0          2          1          1          2          0          0          0
       0
12.101.0.1              bar                       0          0          5          1
         0          4          2          2          4          0          0          0
       0
12.0.0.1                default                   0          0          7          1
         0          6          4          2          4          0          0          0
       0
```

The **show bgp bmp server <ID>** command displays the advertisement-interval time that you configured.

```
Router#show bgp bmp server 1
BMP server 1
Host 12.1.2.1 Port 16001
Connected for 04:22:39
Last Disconnect event received : 00:00:00
Precedence:  internet
BGP neighbors: 10
VRF: - (0x60000000)
Update Source: - (-)
Update Source Vrf ID: 0x0
Update Mode                      : In-Post-Policy, Local-RIB
  In-Post-Policy
   Advertisement interval        : 15 secs
   Scanner interval              : 60 secs
  Local-RIB
   Advertisement interval        : 15 secs
   Scanner interval
```

```
     Global                        : 60 secs
     IPv4 Unicast                  : 60 secs
     VPNv4 Unicast                 : 60 secs
     IPv6 Unicast                  : 60 secs
     VPNv6 Unicast                 : 60 secs
Flapping Delay                     : 300 secs
Initial Delay                      : 0 secs
Initial Refresh Delay              : 1 secs
Initial Refresh Spread             : 1 secs
Stats Reporting Period             : 0 secs
Queue Route Mon Msg buffer limit   : 133693 KB (Current Server Up Count: 2)
Queue Route Mon Msg buffer usage   : 0 B
Queue write pulse sent             : Dec 16 15:19:36.755, not set (all)
-----More------
```

The **show bgp bmp server <ID>** command displays the BMP scan time that you configured.

```
Router#show bgp bmp server 1
BMP server 1
Host 12.1.2.1 Port 16001
Connected for 04:22:39
Last Disconnect event received : 00:00:00
Precedence:  internet
BGP neighbors: 10
VRF: - (0x60000000)
Update Source: - (-)
Update Source Vrf ID: 0x0
Update Mode                        : In-Post-Policy, Local-RIB
  In-Post-Policy
   Advertisement interval          : 15 secs
   Scanner interval                : 60 secs
  Local-RIB
   Advertisement interval          : 15 secs
   Scanner interval
     Global                        : 60 secs
     IPv4 Unicast                  : 60 secs
     VPNv4 Unicast                 : 60 secs
     IPv6 Unicast                  : 60 secs
     VPNv6 Unicast                 : 60 secs
Flapping Delay                     : 300 secs
Initial Delay                      : 0 secs
Initial Refresh Delay              : 1 secs
Initial Refresh Spread             : 1 secs
Stats Reporting Period             : 0 secs
Queue Route Mon Msg buffer limit   : 133693 KB (Current Server Up Count: 2)
Queue Route Mon Msg buffer usage   : 0 B
Queue write pulse sent             : Dec 16 15:19:36.755, not set (all)
-----More------
```

The **show bgp bmp table** command displays summary information of prefixes and paths learnt from the monitored BGP neighbors.

```
Router#show bgp bmp table
---------------------------------------------
Route monitoring mode : Inbound Post-Policy
---------------------------------------------
Address-Family : IPv4 Unicast
-----------------------------------
BMP prefix table version : 11
Number of prefixes: 5, paths: 9
```

```
Network            Version    Neighbor         Bit Map    Flags  Path ID
100.1.1.3/32       7          23.0.1.3         0x03/0xfc  0x02   1
100.1.2.3/32       11         23:0:2::3        0x03/0xfc  0x02   1
 100.1.2.3/32      11         23.0.2.3         0x03/0xfc  0x02   1
100.1.3.3/32       10         23.0.2.3         0x03/0xfc  0x03   1
 100.1.3.3/32      10         23.0.1.3         0x03/0xfc  0x03   1
200.1.1.1/32       5          12::1            0x03/0xfc  0x00
 200.1.1.1/32      5          12.0.0.1         0x03/0xfc  0x00
200.2.1.1/32       6          12::1            0x03/0xfc  0x00
 200.2.1.1/32      6          12.0.0.1         0x03/0xfc  0x00
```

If a prefix has more than one path, the prefix is printed multiple times, once for each path, except that for the second paths onwards, the prefix is indented by one space as visible in case of the prefix 100.1.2.3 in the above table. This prefix has a path via neighbor 23:0:2::3 and 23.0.2.3.

The *Bit Map* displayed above has 2 bitmap values. The first value is the bitmap of BMP servers to which the prefix is advertised and the second value is the bitmap of BMP servers to which the prefix is yet to be advertized. Bit **n** in both the bitmaps (where n is 0 to 7) represent Server **m** where **m** = **n** + 1. For example, bit 3 represents Server 4.

If the *advertized* bitmap is 0x03 (00000011b) it means that the prefix is advertized to servers 1 and 2. Simularly, if the *yet to be advertized* bitmap is 0xfc, then it means that the prefix is yet to be advertized to servers 3, 4, 5, 6, 7 and 8. However if any of these servers is not configured, then the corresponding bits can be ignored.

You can run the **show bgp bmp table** command as follows:

- For a specific address family identifiers (AFI) or sub-address family identifiers (SAFI). For example, **show bgp bmp vpnv4 unicast** table, or a specific BMP monitoring mode, for example **show bgp bmp mode inbound post-policy** table.
- If you do not specify the AFI and SAFI, the command displays information for the default AFI `IPv4 Unicast`.
- If you do not specify the BMP monitoring mode, the command displays information for the default mode `Inbound Post-Policy`.

  For example, if you run the command **show bgp bmp table** where the BMP monitoring mode or the AFI are not specified, the command displays information for `Inbound Post-Policy` BMP mode `IPv4 Unicast` AFI.
- For all BMP monitoring modes (mode all) and/or all AFIs (all all).

The following example shows a brief summary of the prefixes for BMP mode **Local-RIB** and **IPv4 Unicast** AFI:

```
Router#show bgp bmp mode local-rib table
-------------------------------------------
Route monitoring mode : Local-RIB
-------------------------------------------
Address-Family : IPv4 Unicast
-----------------------------------
BGP main routing table version : 9
Number of prefixes: 4, paths: 4

   Network            Next Hop            Metric LocPrf Path
   100.1.1.3/32       23.0.1.3                 0    100 ?
   100.1.2.3/32       23.0.2.3                 0    100 ?
   200.1.1.1/32       12.0.0.1                 0        1 ?
   200.2.1.1/32       12.0.0.1                 0        1 ?
```

```
Processed 4 prefixes, 4 paths
```

The following example shows a brief summary of the prefixes for BMP mode **Inbound Post-Policy** and **VPNv4 Unicast** AFI:

```
Router#show bgp bmp vpnv4 unicast table

-----------------------------------------------
Route monitoring mode : Inbound Post-Policy
-----------------------------------------------
Address-Family : VPNv4 Unicast
-----------------------------------
BMP prefix table version : 14
Number of RDs: 3, prefixes: 6, paths: 10
   Network            Version    Neighbor         Bit Map    Flags  Path ID
Route Distinguisher: 2:100 (default for vrf foo), Version 6
   200.1.1.1/32       5          12.100.0.1       0x03/0xfc  0x00
   200.2.1.1/32       6          12.100.0.1       0x03/0xfc  0x00
Route Distinguisher: 2:101 (default for vrf bar), Version 8
   200.1.1.1/32       7          12.101.0.1       0x03/0xfc  0x00
   200.2.1.1/32       8          12.101.0.1       0x03/0xfc  0x00
Route Distinguisher: 2:200, Version 14
   200.1.1.4/32       14         23.0.1.3         0x03/0xfc  0x02
    200.1.1.4/32      14         23.0.2.3         0x03/0xfc  0x02
    200.1.1.4/32      14         23:0:2::3        0x03/0xfc  0x02
   200.2.1.4/32       13         23.0.1.3         0x03/0xfc  0x02
    200.2.1.4/32      13         23.0.2.3         0x03/0xfc  0x02
    200.2.1.4/32      13         23:0:2::3        0x03/0xfc  0x02

Processed 6 prefixes, 10 paths
```

The following example shows a brief summary of the prefixes for **all** BMP modes and **all** AFI:

```
Router#show bgp bmp mode all all all table

-----------------------------------------------
Route monitoring mode : Inbound Post-Policy
-----------------------------------------------
Address-Family : VPNv4 Unicast
-----------------------------------
BMP prefix table version : 14
Number of RDs: 3, prefixes: 6, paths: 10

   Network            Version    Neighbor         Bit Map    Flags  Path ID
Route Distinguisher: 2:100 (default for vrf foo), Version 6
   200.1.1.1/32       5          12.100.0.1       0x03/0xfc  0x00
   200.2.1.1/32       6          12.100.0.1       0x03/0xfc  0x00
Route Distinguisher: 2:101 (default for vrf bar), Version 8
   200.1.1.1/32       7          12.101.0.1       0x03/0xfc  0x00
   200.2.1.1/32       8          12.101.0.1       0x03/0xfc  0x00
Route Distinguisher: 2:200, Version 14
   200.1.1.4/32       14         23.0.1.3         0x03/0xfc  0x02
    200.1.1.4/32      14         23.0.2.3         0x03/0xfc  0x02
    200.1.1.4/32      14         23:0:2::3        0x03/0xfc  0x02
   200.2.1.4/32       13         23.0.1.3         0x03/0xfc  0x02
    200.2.1.4/32      13         23.0.2.3         0x03/0xfc  0x02
    200.2.1.4/32      13         23:0:2::3        0x03/0xfc  0x02

Processed 6 prefixes, 10 paths
Address-Family : VPNv6 Unicast
-----------------------------------
```

```
BMP prefix table version : 16
Number of RDs: 3, prefixes: 6, paths: 12

   Network            Version    Neighbor         Bit Map    Flags  Path ID
Route Distinguisher: 2:100 (default for vrf foo), Version 8
   200:1:1::1/128     7          12.100.0.1       0x03/0xfc  0x00
    200:1:1::1/128    7          12:100::1        0x03/0xfc  0x00
   200:2:1::1/128     8          12.100.0.1       0x03/0xfc  0x00
    200:2:1::1/128    8          12:100::1        0x03/0xfc  0x00

Route Distinguisher: 2:101 (default for vrf bar), Version 10
   200:1:1::1/128     9          12.101.0.1       0x03/0xfc  0x00
   200:2:1::1/128     10         12.101.0.1       0x03/0xfc  0x00

Route Distinguisher: 2:200, Version 16
   200:1:1::4/128     16         23.0.2.3         0x03/0xfc  0x02
    200:1:1::4/128    16         23:0:2::3        0x03/0xfc  0x02
    200:1:1::4/128    16         23.0.1.3         0x03/0xfc  0x02
   200:2:1::4/128     15         23.0.2.3         0x03/0xfc  0x02
    200:2:1::4/128    15         23:0:2::3        0x03/0xfc  0x02
    200:2:1::4/128    15         23.0.1.3         0x03/0xfc  0x02

Processed 6 prefixes, 12 paths
Address-Family : IPv4 Labeled-unicast
------------------------------------
BMP prefix table version : 11
Number of prefixes: 5, paths: 9
   Network            Version    Neighbor         Bit Map    Flags  Path ID
   100.1.3.3/32       10         23.0.2.3         0x03/0xfc  0x03   1
    100.1.3.3/32      10         23.0.1.3         0x03/0xfc  0x03   1

Processed 1 prefixes, 2 paths
Address-Family : IPv6 Labeled-unicast
------------------------------------
BMP prefix table version : 12
Number of prefixes: 5, paths: 10

   Network            Version    Neighbor         Bit Map    Flags  Path ID
   100:1:3::3/128     12         23:0:2::3        0x03/0xfc  0x03   1
    100:1:3::3/128    12         23.0.2.3         0x03/0xfc  0x03   1
    100:1:3::3/128    12         23.0.1.3         0x03/0xfc  0x03   1

Processed 1 prefixes, 3 paths
Address-Family : IPv4 Unicast
------------------------------------
BMP prefix table version : 11
Number of prefixes: 5, paths: 9

   Network            Version    Neighbor         Bit Map    Flags  Path ID
   100.1.1.3/32       7          23.0.1.3         0x03/0xfc  0x02   1
   100.1.2.3/32       11         23:0:2::3        0x03/0xfc  0x02   1
    100.1.2.3/32      11         23.0.2.3         0x03/0xfc  0x02   1
   100.1.3.3/32       10         23.0.2.3         0x03/0xfc  0x03   1
    100.1.3.3/32      10         23.0.1.3         0x03/0xfc  0x03   1
   200.1.1.1/32       5          12::1            0x03/0xfc  0x00
    200.1.1.1/32      5          12.0.0.1         0x03/0xfc  0x00
   200.2.1.1/32       6          12::1            0x03/0xfc  0x00
    200.2.1.1/32      6          12.0.0.1         0x03/0xfc  0x00

Processed 5 prefixes, 9 paths
Address-Family : IPv6 Unicast
------------------------------------
BMP prefix table version : 12
Number of prefixes: 5, paths: 10
```

```
   Network            Version    Neighbor         Bit Map    Flags  Path ID
   100:1:1::3/128     8          23.0.1.3         0x03/0xfc  0x02   1
   100:1:2::3/128     9          23:0:2::3        0x03/0xfc  0x02   1
    100:1:2::3/128    9          23.0.2.3         0x03/0xfc  0x02   1
   100:1:3::3/128     12         23:0:2::3        0x03/0xfc  0x03   1
    100:1:3::3/128    12         23.0.2.3         0x03/0xfc  0x03   1
    100:1:3::3/128    12         23.0.1.3         0x03/0xfc  0x03   1
   200:1:1::1/128     5          12.0.0.1         0x03/0xfc  0x00
    200:1:1::1/128    5          12::1            0x03/0xfc  0x00
   200:2:1::1/128     6          12.0.0.1         0x03/0xfc  0x00
    200:2:1::1/128    6          12::1            0x03/0xfc  0x00

Processed 5 prefixes, 10 paths
-----------------------------------------------
Route monitoring mode : Local-RIB
-----------------------------------------------
Address-Family : VPNv4 Unicast
-----------------------------------
BGP main routing table version : 33
Number of prefixes: 10, paths: 10

   Network            Next Hop            Metric LocPrf Path
Route Distinguisher: 2:100 (default for vrf foo), Version 33
   200.1.1.1/32       12.100.0.1               0        1 ?
   200.1.1.4/32       23.0.1.3                 0    100 3 ?
   200.2.1.1/32       12.100.0.1               0        1 ?
   200.2.1.4/32       23.0.1.3                 0    100 3 ?

Route Distinguisher: 2:101 (default for vrf bar), Version 31
   200.1.1.1/32       12.101.0.1               0        1 ?
   200.1.1.4/32       23.0.1.3                 0    100 3 ?
   200.2.1.1/32       12.101.0.1               0        1 ?
   200.2.1.4/32       23.0.1.3                 0    100 3 ?

Route Distinguisher: 2:200, Version 27
   200.1.1.4/32       23.0.1.3                 0    100 3 ?
   200.2.1.4/32       23.0.1.3                 0    100 3 ?
Processed 10 prefixes, 10 paths

Address-Family : VPNv6 Unicast
-----------------------------------
BGP main routing table version : 31

Number of prefixes: 10, paths: 10

   Network            Next Hop            Metric LocPrf Path
Route Distinguisher: 2:100 (default for vrf foo), Version 29
   200:1:1::1/128     12:100::1                0        1 ?
   200:1:1::4/128     23.0.1.3                 0    100 3 ?
   200:2:1::1/128     12:100::1                0        1 ?
   200:2:1::4/128     23.0.1.3                 0    100 3 ?

Route Distinguisher: 2:101 (default for vrf bar), Version 23
   200:1:1::1/128     12:101::1                0        1 ?
   200:1:1::4/128     23.0.1.3                 0    100 3 ?
   200:2:1::1/128     12:101::1                0        1 ?
   200:2:1::4/128     23.0.1.3                 0    100 3 ?

Route Distinguisher: 2:200, Version 31

   200:1:1::4/128     23.0.1.3                 0    100 3 ?
   200:2:1::4/128     23.0.1.3                 0    100 3 ?
```

```
Processed 10 prefixes, 10 paths
Address-Family : IPv4 Labeled-unicast
-------------------------------------
BGP main routing table version : 9
Number of prefixes: 1, paths: 1

   Network            Next Hop            Metric LocPrf Path

   100.1.3.3/32       23.0.1.3                 0    100 ?

Processed 1 prefixes, 1 paths
Address-Family : IPv6 Labeled-unicast
-------------------------------------
BGP main routing table version : 9
Number of prefixes: 1, paths: 1

   Network            Next Hop            Metric LocPrf Path

   100:1:3::3/128     23.0.1.3                 0    100 ?
Processed 1 prefixes, 1 paths
Address-Family : IPv4 Unicast
-------------------------------------
BGP main routing table version : 9

Number of prefixes: 4, paths: 4

   Network            Next Hop            Metric LocPrf Path
   100.1.1.3/32       23.0.1.3                 0    100 ?
   100.1.2.3/32       23.0.2.3                 0    100 ?
   200.1.1.1/32       12.0.0.1                 0        1 ?
   200.2.1.1/32       12.0.0.1                 0        1 ?

Processed 4 prefixes, 4 paths
Address-Family : IPv6 Unicast
-------------------------------------
BGP main routing table version : 9
Number of prefixes: 4, paths: 4

   Network            Next Hop            Metric LocPrf Path
   100:1:1::3/128     23:0:1::3                0    100 ?
   100:1:2::3/128     23:0:2::3                0    100 ?
   200:1:1::1/128     12::1                    0        1 ?
   200:2:1::1/128     12::1                    0        1 ?

Processed 4 prefixes, 4 paths
```

Note Specifying AFI and SAFI after the **table** keyword option is not supported from Release 7.5.4. So, you can specify the same before the **table** keyword as follows:

```
show bgp bmp [mode <bmp-mode>] [<afi> <safi>] table
```

The option to specify AFIs and SAFIs after the **table** keyword is completely removed from Cisco IOS XR Release 24.1.1.

The tables for VPN AFI can be filtered on a Route Distinguisher (RD) or a VRF. The following example shows the table for **VPNv4 Unicast** without RD or VRF filter.

```
Router#show bgp bmp vpnv4 unicast table
```

```
-------------------------------------------
Route monitoring mode : Inbound Post-Policy
-------------------------------------------
Address-Family : VPNv4 Unicast
-----------------------------------
BMP prefix table version : 18

Number of RDs: 3, prefixes: 6, paths: 12

   Network            Version    Neighbor         Bit Map    Flags  Path ID
Route Distinguisher: 2:100 (default for vrf foo), Version 14
   200.1.1.1/32       13         12.100.0.1       0x03/0xfc  0x02
   200.2.1.1/32       14         12.100.0.1       0x03/0xfc  0x02

Route Distinguisher: 2:101 (default for vrf bar), Version 16
   200.1.1.1/32       15         12.101.0.1       0x03/0xfc  0x02
   200.2.1.1/32       16         12.101.0.1       0x03/0xfc  0x02

Route Distinguisher: 2:200, Version 18
   200.1.1.4/32       17         23.0.3.3         0x03/0xfc  0x00
    200.1.1.4/32      17         23.0.2.3         0x03/0xfc  0x00
    200.1.1.4/32      17         23.0.4.3         0x03/0xfc  0x00
    200.1.1.4/32      17         23.0.1.3         0x03/0xfc  0x00
   200.2.1.4/32       18         23.0.3.3         0x03/0xfc  0x00
    200.2.1.4/32      18         23.0.2.3         0x03/0xfc  0x00
    200.2.1.4/32      18         23.0.4.3         0x03/0xfc  0x00
    200.2.1.4/32      18         23.0.1.3         0x03/0xfc  0x00
```

The following example shows the **show bgp bmp vpnv4 unicast table** filtered on route distinguisher *2:100*:

```
Router#show bgp bmp vpnv4 unicast table rd 2:100

-------------------------------------------
Route monitoring mode : Inbound Post-Policy
-------------------------------------------
Address-Family : VPNv4 Unicast
-----------------------------------

BMP prefix table version : 18
Number of RDs: 3, prefixes: 6, paths: 12

   Network            Version    Neighbor         Bit Map    Flags  Path ID
Route Distinguisher: 2:100 (default for vrf foo), Version 14
   200.1.1.1/32       13         12.100.0.1       0x03/0xfc  0x02
   200.2.1.1/32       14         12.100.0.1       0x03/0xfc  0x02

Processed 2 prefixes, 2 paths
```

The following example shows the **show bgp bmp vpnv4 unicast table** filtered on VRF *foo*:

```
Router#show bgp bmp vpnv4 unicast table vrf foo
-------------------------------------------
Route monitoring mode : Inbound Post-Policy
-------------------------------------------
Address-Family : VPNv4 Unicast
-----------------------------------
BMP prefix table version : 18

Number of RDs: 3, prefixes: 6, paths: 12

   Network            Version    Neighbor         Bit Map    Flags  Path ID
Route Distinguisher: 2:100 (default for vrf foo), Version 14
```

```
   200.1.1.1/32       13         12.100.0.1       0x03/0xfc  0x02
   200.2.1.1/32       14         12.100.0.1       0x03/0xfc  0x02

Processed 2 prefixes, 2 paths
```

When the **show bgp bmp table** command is run with prefix filter, it displays the details of the specified prefix as shown in the following example:

```
Router#show bgp bmp table 100.1.1.3/32
-------------------------------------------
Route monitoring mode : Inbound Post-Policy
-------------------------------------------
BGP routing table entry for 100.1.1.3/32

Versions:
  Process           Net
  Speaker           7
Last Modified: Dec 16 15:05:22.000 for 07:05:14
    Flags: 0x00000000+0x00000000
    Net Ptr: 0x7f0bbd65ff34
Paths: (1 available)
  Path #1: Received by speaker 0
  Flags: 0x00000000
  Local
  from 23.0.1.3
  Nexthop: 23.0.1.3
  Received Path ID: 1
  Received Local Preference: 100
  Path Advertised to BMP server(s) - (Map: 0x03): Server ID: 1, 2
  Path pending send to BMP server(s) - (Map: 0xfc): Server ID: None
  Path Ptr : 0x7f0bb806f0d0, BGP Path Ptr : 0x7f0bbd915f68
```

The following examples displays the output of the **show bgp bmp table** command with the **summary** keyword. It displays the summary of the table of the specified AFI or SAFI, or the default AFI or SAFI if not specified in the command.

```
Router#show bgp bmp table summary
-------------------------------------------
Route monitoring mode : Inbound Post-Policy
-------------------------------------------
Address-Family : IPv4 Unicast
-----------------------------------
BMP prefix table version : 12

Number of prefixes: 6, paths: 6
```

The following examples displays the output when you run the **show bgp bmp table** command with the **neighbor** filter. It displays the prefixes leanrt from the specified neighbor.

```
Router#show bgp bmp table neighbor 23.0.2.3

-------------------------------------------
Route monitoring mode : Inbound Post-Policy
-------------------------------------------
Address-Family : IPv4 Unicast
-----------------------------------
BMP prefix table version : 11
```

```
Number of prefixes: 5, paths: 9

   Network            Version    Neighbor         Bit Map    Flags  Path ID
   100.1.2.3/32       11         23.0.2.3         0x03/0xfc  0x02   1
   100.1.3.3/32       10         23.0.2.3         0x03/0xfc  0x03   1
```

The following example displays filtering of prefixes based on prefix version range. Only the prefixes with version within the specified range are displayed.

```
Router#show bgp bmp table version begin 7 end 15
-----------------------------------------------
Route monitoring mode : Inbound Post-Policy
-----------------------------------------------
Address-Family : IPv4 Unicast
-----------------------------------
BMP prefix table version : 11

Number of prefixes: 5, paths: 9
   Network            Version    Neighbor         Bit Map    Flags  Path ID
   100.1.1.3/32       7          23.0.1.3         0x03/0xfc  0x02   1
   100.1.2.3/32       11         23:0:2::3        0x03/0xfc  0x02   1
    100.1.2.3/32      11         23.0.2.3         0x03/0xfc  0x02   1
   100.1.3.3/32       10         23.0.2.3         0x03/0xfc  0x03   1
    100.1.3.3/32      10         23.0.1.3         0x03/0xfc  0x03   1

Processed 3 prefixes, 5 paths
```

The command **show bgp bmp convergence** displays whether there is any pending work for BMP to perform and whether a specific address family has reached convergence. The command is applicable to **Inbound Post-Policy** as well as **Local-RIB** modes. Convergence is declared when the following conditions are met.

- Table copy is not pending (applicable only for **Inbound Post-Policy**)
- Update generation is not pending
- All BMP servers that are not administratively shutdown are *Up* state
- All BMP servers have caught up to the table version
- Write queues of all BMP servers are empty

```
Router#show bgp bmp convergence
-----------------------------------------------
Route monitoring mode : Inbound Post-Policy
-----------------------------------------------
Address-Family : IPv4 Unicast
-----------------------------------
Converged.

All paths from monitored neighbors are sent to the BMP servers that are not administratively
 shutdown.
BMP write queue is empty.
```

The following example displays an address family that is converged. Since BMP mode is not specified, convergence details are shown for the default mode inbound post-policy.

```
Router#show bgp bmp vpnv4 unicast convergence
```

```
-------------------------------------------
Route monitoring mode : Inbound Post-Policy
-------------------------------------------
Address-Family : VPNv4 Unicast
-----------------------------------

Converged.

All paths from monitored neighbors are sent to the BMP servers that are not administratively
 shutdown.
BMP write queue is empty.
```

The following example displays the convergence details for all AFIs and SAFIs. Since BMP mode is not specified, convergence information is shown for the default mode inbound post-policy.

```
Router#show bgp bmp all all convergence
-------------------------------------------

Route monitoring mode : Inbound Post-Policy
-------------------------------------------
Address-Family : VPNv4 Unicast
-----------------------------------
Converged.

All paths from monitored neighbors are sent to the BMP servers that are not administratively
 shutdown.

BMP write queue is empty.

Address-Family : VPNv6 Unicast
-----------------------------------
Converged.

All paths from monitored neighbors are sent to the BMP servers that are not administratively
 shutdown.

BMP write queue is empty.

Address-Family : IPv4 Labeled-unicast
-----------------------------------
Converged.

All paths from monitored neighbors are sent to the BMP servers that are not administratively
 shutdown.

BMP write queue is empty.

Address-Family : IPv6 Labeled-unicast
-----------------------------------

Converged.

All paths from monitored neighbors are sent to the BMP servers that are not administratively
 shutdown.

BMP write queue is empty.

Address-Family : IPv4 Unicast
-----------------------------------
Converged.

All paths from monitored neighbors are sent to the BMP servers that are not administratively
 shutdown.
```

```
BMP write queue is empty.

Address-Family : IPv6 Unicast
-----------------------------------
Converged.

All paths from monitored neighbors are sent to the BMP servers that are not administratively
 shutdown.

BMP write queue is empty.
```

The following example displays the convergence details for BMP mode **inbound post-policy** explicitly specified. The output of this command is same as that of **show bgp bmp mode convergence** as the explicitly specified mode is same as the default mode.

```
Router#show bgp bmp mode inbound post-policy convergence
-----------------------------------------------
Route monitoring mode : Inbound Post-Policy
-----------------------------------------------
Address-Family : IPv4 Unicast
-----------------------------------

Converged.

All paths from monitored neighbors are sent to the BMP servers that are not administratively
 shutdown.
BMP write queue is empty.
```

The following examples displays the convergence details for BMP **Local-RIB**.

```
Router#show bgp bmp mode local-rib convergence
-----------------------------------------------
Route monitoring mode : Local-RIB
-----------------------------------------------
Address-Family : IPv4 Unicast
-----------------------------------

Converged.

All paths from monitored neighbors are sent to the BMP servers that are not administratively
 shutdown.
BMP write queue is empty.
```

The following examples displays the convergence details for non-default address family **VPNv4 Unicast**.

```
Router#show bgp bmp mode local-rib vpnv4 unicast convergence

-----------------------------------------------
Route monitoring mode : Local-RIB
-----------------------------------------------
Address-Family : VPNv4 Unicast
-----------------------------------

Converged.

All paths from monitored neighbors are sent to the BMP servers that are not administratively
 shutdown.
BMP write queue is empty.
```

The following examples displays the convergence details for all address families for BMP mode **Local-RIB**.

```
Router#show bgp bmp mode local-rib all all convergence
-----------------------------------------------
Route monitoring mode : Local-RIB
-----------------------------------------------
Address-Family : VPNv4 Unicast
-----------------------------------
Converged.

All paths from monitored neighbors are sent to the BMP servers that are not administratively
 shutdown.
BMP write queue is empty.

Address-Family : VPNv6 Unicast
-----------------------------------
Converged.

All paths from monitored neighbors are sent to the BMP servers that are not administratively
 shutdown.
BMP write queue is empty.

Address-Family : IPv4 Labeled-unicast
-----------------------------------
Converged.

All paths from monitored neighbors are sent to the BMP servers that are not administratively
 shutdown.

BMP write queue is empty.

Address-Family : IPv6 Labeled-unicast
-----------------------------------
Converged.

All paths from monitored neighbors are sent to the BMP servers that are not administratively
 shutdown.
BMP write queue is empty.

Address-Family : IPv4 Unicast
-----------------------------------
Converged.

All paths from monitored neighbors are sent to the BMP servers that are not administratively
 shutdown.
BMP write queue is empty.

Address-Family : IPv6 Unicast
-----------------------------------
Converged.

All paths from monitored neighbors are sent to the BMP servers that are not administratively
 shutdown.
BMP write queue is empty.
```

The following example shows the convergence details of an address family that has not converged (in this case the default address family **ipv4 unicast** since it is not specified in the command). In this example, in case of **Inbound Post-Policy**, convergence is not reached as the servers are in Down state and they have not caught up to the **Table** version:

```
Router#show bgp bmp mode all convergence
-----------------------------------------------
Route monitoring mode : Inbound Post-Policy
-----------------------------------------------
Address-Family : IPv4 Unicast
-----------------------------------
Not converged.
All paths from monitored neighbors may not be sent to the BMP servers that are not
administratively shutdown.

Table copy pending: N
Update generation pending: N
Table version: 108

Server  Caught-Up Version  Queue Size  State
-----------------------------------------------
      1                  0           0  Down

      2                  0           0  Down
-----------------------------------------------
Route monitoring mode : Local-RIB
-----------------------------------------------
Address-Family : IPv4 Unicast
-----------------------------------
Not converged.

All paths from monitored neighbors may not be sent to the BMP servers that are not
administratively shutdown.

Table copy pending: N
Update generation pending: N
Table version: 28

Server  Caught-Up Version  Queue Size  State
-----------------------------------------------
      1                  0           0  Down

      2                  0           0  Down
```

In the following example of the **show bgp bmp mode all convergence** command, convergence is reached for BMP mode local-rib and address-family ipv4 unicast. However, the same is not yet reached for BMP mode inbound post-policy and address-family ipv4 unicast.

```
Router#show bgp bmp mode all convergence
-----------------------------------------------
Route monitoring mode : Inbound Post-Policy
-----------------------------------------------
Address-Family : IPv4 Unicast
-----------------------------------
Not converged.

All paths from monitored neighbors may not be sent to the BMP servers that are not
administratively shutdown.

Table copy pending: N
Update generation pending: Y
Table version: 108

 Server  Caught-Up Version  Queue Size  State
-----------------------------------------------
      1                  0           0  Up
```

```
      2                  0           0  Up
-----------------------------------------
Route monitoring mode : Local-RIB
-----------------------------------------
Address-Family : IPv4 Unicast
-----------------------------------

Converged.

All paths from monitored neighbors are sent to the BMP servers that are not administratively
 shutdown.
BMP write queue is empty.
```

In the following example of the same command, convergence is still not reached for BMP mode inbound post-policy. In this case the server are in *Up* state and they have caught up to the **Table** version. However, the queue for Server 1 is not empty.

```
Router#show bgp bmp mode all convergence
-----------------------------------------
Route monitoring mode : Inbound Post-Policy
-----------------------------------------
Address-Family : IPv4 Unicast
-----------------------------------
Not converged.

All paths from monitored neighbors may not be sent to the BMP servers that are not
administratively shutdown.

Table copy pending: N
Update generation pending: N
Table version: 108

Server  Caught-Up Version  Queue Size  State
-------------------------------------------------
      1                108         432  Up

      2                108           0  Up
-----------------------------------------
Route monitoring mode : Local-RIB
-----------------------------------------
Address-Family : IPv4 Unicast
-----------------------------------

Converged.

All paths from monitored neighbors are sent to the BMP servers that are not administratively
 shutdown.
BMP write queue is empty.
```

Note The origin of prefixes conveyed in the **Inbound Post-Policy** view conveys is the actual BGP peers. Whereas the origin of prefixes conveyed in the **Local-RIB** view is the **Local-RIB** itself. The **Local-RIB** is viewed as a set of **emulated** peers from which the prefixes originate. Each emulated peer is associated with the prefixes belonging to one VRF (all address families inclusive) that are conveyed to one BMP server. In other words, one emulated peer is per-VRF per-BMP server. So, if BGP has *n* VRFs and there are *m* BMP servers, then *n* x *m* emulated peers are created automatically.

The internal construct of the emulated peer is largely same as that of the normal BGP peer. So, the output of some of the emulated peers related commands resembles that of the BGP peer related commands.

The following example shows the details of the Local RIB emulated peers.

Note This keyword **emulated-peers** is applicable only to BMP mode **Local RIB**.

```
Router#show bgp bmp mode local-rib emulated-peers
BGP emulated peer is 0.0.0.1
Remote AS 2, local AS 2, internal link
Remote router ID 100.1.1.2
  BGP state = Established, up for 08:05:38
  Previous State: Idle
  Last Received Message: None
  BGP Monitoring(BMP) activated for servers:
      1
  Multi-protocol capability received
  Emulated peer capabilities:
    Route refresh: advertised (old + new)
    4-byte AS: advertised and received
    Address family IPv4 Unicast: advertised and received
    Address family IPv4 Labeled-unicast: advertised and received
    Address family VPNv4 Unicast: advertised and received
    Address family IPv6 Unicast: advertised and received
    Address family IPv6 Labeled-unicast: advertised and received
    Address family VPNv6 Unicast: advertised and received
  Received 0 messages, 0 notifications, 0 in queue
  Sent 20 messages, 0 notifications, 0 in queue
  Minimum time between advertisement runs is 15 secs

For Address Family: IPv4 Unicast
  BMP emulated peer version 9
  Update group: 0.9 Filter-group: 0.8  No Refresh request being processed
  Route-Reflector Client
    Extended Nexthop Encoding: advertised and received
  Route refresh request: received 0, sent 0
  Prefix advertised 4, suppressed 0, withdrawn 0
  AIGP is enabled
  Last ack version 9, Last synced ack version 0
  Outstanding version objects: current 0, max 1, refresh 0
  Additional-paths operation: None
  Send Multicast Attributes
  Advertise routes with local-label via Unicast SAFI

For Address Family: IPv4 Labeled-unicast
  BMP emulated peer version 9
```

```
  Update group: 0.10 Filter-group: 0.9  No Refresh request being processed
  Route-Reflector Client
    Extended Nexthop Encoding: advertised and received
  Route refresh request: received 0, sent 0
  Prefix advertised 1, suppressed 0, withdrawn 0
  AIGP is enabled
  Last ack version 9, Last synced ack version 0
  Outstanding version objects: current 0, max 1, refresh 0
  Additional-paths operation: None
  Send Multicast Attributes

For Address Family: VPNv4 Unicast
  BMP emulated peer version 33
  Update group: 0.7 Filter-group: 0.6  No Refresh request being processed
  Route-Reflector Client
    Extended Nexthop Encoding: advertised and received
  Route refresh request: received 0, sent 0
  Prefix advertised 6, suppressed 0, withdrawn 0
  AIGP is enabled
  Last ack version 33, Last synced ack version 0
  Outstanding version objects: current 0, max 1, refresh 0
  Additional-paths operation: None
  Send Multicast Attributes

For Address Family: IPv6 Unicast
  BMP emulated peer version 9
  Update group: 0.11 Filter-group: 0.9  No Refresh request being processed
  Route-Reflector Client
    Extended Nexthop Encoding: advertised and received
  Route refresh request: received 0, sent 0
  Prefix advertised 4, suppressed 0, withdrawn 0
  AIGP is enabled
  Last ack version 9, Last synced ack version 0
  Outstanding version objects: current 0, max 1, refresh 0
  Additional-paths operation: None
  Send Multicast Attributes
  Advertise routes with local-label via Unicast SAFI

For Address Family: IPv6 Labeled-unicast
  BMP emulated peer version 9
  Update group: 0.12 Filter-group: 0.10  No Refresh request being processed
  Route-Reflector Client
    Extended Nexthop Encoding: advertised and received
  Route refresh request: received 0, sent 0
  Prefix advertised 1, suppressed 0, withdrawn 0
  AIGP is enabled
  Last ack version 9, Last synced ack version 0
  Outstanding version objects: current 0, max 1, refresh 0
  Additional-paths operation: None
  Send Multicast Attributes

For Address Family: VPNv6 Unicast
  BMP emulated peer version 31
  Update group: 0.7 Filter-group: 0.6  No Refresh request being processed
  Route-Reflector Client
    Extended Nexthop Encoding: advertised and received
  Route refresh request: received 0, sent 0
  Prefix advertised 6, suppressed 0, withdrawn 0
  AIGP is enabled
  Last ack version 31, Last synced ack version 0
  Outstanding version objects: current 0, max 1, refresh 0
  Additional-paths operation: None
  Send Multicast Attributes
```

```
  Connections established 1; dropped 0
  Last reset 00:00:00
```

The **show bgp bmp mode local-rib emulated peers** command with additional keyword **brief** is used to display a brief summary of the Local RIB emulated peers.

```
Router#show bgp bmp mode local-rib emulated-peers brief

EmuPeerId       BMPServer  VRF                    Spk    AS  OutQ  State
0.0.0.1                 1  default                  0     2     0  Established
0.0.0.2                 2  default                  0     2     0  Established
0.0.0.1                 1  bar                      0     2     0  Established
0.0.0.2                 2  bar                      0     2     0  Established
0.0.0.1                 1  foo                      0     2     0  Established
0.0.0.2                 2  foo                      0     2     0  Established
```

The following example shows the detailed summary about Local RIB emulated peers when the **detail** keyword is used.

```
Router#show bgp bmp mode local-rib emulated-peers detail

BGP emulated peer is 0.0.0.1
 Remote AS 2, local AS 2, internal link
 Remote router ID 100.1.1.2
  BGP state = Established, up for 2d20h
  Previous State: Idle
  Last Received Message: None
  BGP Monitoring(BMP) activated for servers:
      1
  Multi-protocol capability received
  Emulated peer capabilities:       Adv         Rcvd
    Route refresh:                  Yes         No
    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 IPv6 Labeled-unicast:  Yes         Yes
    Address family VPNv6 Unicast:   Yes         Yes

  Message stats:
    OutQ depth: 0
  Minimum time between advertisement runs is 15 secs
```

The following example shows the detailed summary of a **Local RIB** emulated peer specified by VRF and emulated-peer address.

```
Router#show bgp bmp mode local-rib vrf foo emulated-peers 0.0.0.1 detail

BGP emulated peer is 0.0.0.1, vrf foo
 Remote AS 2, local AS 2, internal link
 Remote router ID 100.1.1.2
  BGP state = Established, up for 2d21h
  Previous State: Idle
  Last Received Message: None
  BGP Monitoring(BMP) activated for servers:
      1
  Multi-protocol capability received
  Emulated peer capabilities:       Adv         Rcvd
    Route refresh:                  Yes         No
    4-byte AS:                      Yes         Yes
    Address family IPv4 Unicast:    Yes         Yes
```

```
    Address family IPv6 Unicast:    Yes         Yes

  Message stats:
    OutQ depth: 0
  Minimum time between advertisement runs is 15 secs
```

The following example shows the performance statistics information of a **Local RIB** emulated peer specified with VRF and emulated-peer address.

```
Router#show bgp bmp mode local-rib vrf foo emulated-peers 0.0.0.1 performance-statistics

BGP neighbor is 0.0.0.1, vrf foo
 Remote AS 2
  Read 0 messages (0 bytes) in 0 calls (time spent: 0.000 secs)
  Read partly throttled 0 times
  Max socket read size: 0 bytes
  Processed 0 inbound update messages (time spent: 0.000 secs)
  Wrote 750 bytes in 2 calls (time spent: 0.000 secs)
  Processing sub-group: wrote 6 messages in 4 calls (time spent: 0.000 secs)
  Processing write queue: wrote 0 messages in 0 calls (time spent: 0.000 secs)

  Received 0 messages, 0 notifications, 0 in queue
  Sent 6 messages, 0 notifications, 0 in queue
  NSR State: None
  Nbr Primary fd:-1 Reset flags:0x0 Sync flags:0x0 Nbr-flags:0x3000c0 Rst retries:0

 For Address Family: IPv4 Unicast
  Message statistics:
    Update messages processed: 3
    Update messages sent: 3
    Split-horizon Advertisement messages: sent 0, blocked 0
    Split-horizon Withdraw messages: sent 0, blocked 0

 For Address Family: IPv6 Unicast
  Message statistics:
    Update messages processed: 3
    Update messages sent: 3
    Split-horizon Advertisement messages: sent 0, blocked 0
    Split-horizon Withdraw messages: sent 0, blocked 0
```

The following example displays the details of all emulated peers belonging to all VRFs.

```
Router#show bgp bmp mode local-rib vrf all emulated-peers

VRF: bar
--------

BGP emulated peer is 0.0.0.1, vrf bar
 Remote AS 2, local AS 2, internal link
 Remote router ID 100.1.1.2
  BGP state = Established, up for 2d22h
  Previous State: Idle
  Last Received Message: None
  BGP Monitoring(BMP) activated for servers:
      1
  Multi-protocol capability received
  Emulated peer capabilities:
    Route refresh: advertised (old + new)
    4-byte AS: advertised and received
    Address family IPv4 Unicast: advertised and received
    Address family IPv6 Unicast: advertised and received
  Received 0 messages, 0 notifications, 0 in queue
  Sent 6 messages, 0 notifications, 0 in queue
```

```
  Minimum time between advertisement runs is 15 secs

 For Address Family: IPv4 Unicast

  BMP emulated peer version 25
  Update group: 0.1 Filter-group: 0.3  No Refresh request being processed
  Route-Reflector Client
```

The following example displays the details of all emulated peers belong to a specified VRF.

```
Router#show bgp bmp mode local-rib vrf foo emulated-peers

BGP emulated peer is 0.0.0.1, vrf foo
 Remote AS 2, local AS 2, internal link
 Remote router ID 100.1.1.2
  BGP state = Established, up for 2d22h
  Previous State: Idle
  Last Received Message: None
  BGP Monitoring(BMP) activated for servers:
      1
  Multi-protocol capability received
  Emulated peer capabilities:
    Route refresh: advertised (old + new)
    4-byte AS: advertised and received
    Address family IPv4 Unicast: advertised and received
    Address family IPv6 Unicast: advertised and received
  Received 0 messages, 0 notifications, 0 in queue
  Sent 6 messages, 0 notifications, 0 in queue
  Minimum time between advertisement runs is 15 secs

 For Address Family: IPv4 Unicast

  BMP emulated peer version 25
  Update group: 0.1 Filter-group: 0.3  No Refresh request being processed
  Route-Reflector Client
```

Related Commands

Command	Description
bmp server, on page 123	Configures BMP server.
bmp-activate, on page 119	Enables BMP logging for a neighbor.

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 EXECEXEC mode 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

EXECEXEC mode

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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

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 *System Management Command Reference for Cisco ASR 9000 Series Routers* 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 EXECEXEC mode mode:

```
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 10: 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.

Related Commands

Command	Description
route-policy (BGP), on page 290	Applies a routing policy to updates advertised to or received from a BGP neighbor
set default-afi	Sets the default Address Family Identifier (AFI) for the current session.
set default-safi	Sets the default Subaddress Family Identifier (SAFI) for the current session.
show bgp, on page 312	Displays BGP routes.

show bgp community

To display routes that have the specified Border Gateway Protocol (BGP) communities, use the **show bgp community** command in EXECEXEC mode mode.

show bgp [**ipv4** {**unicast** | **multicast** | **labeled-unicast** | **all** | **tunnel**}] **community** *community-list* [**exact-match**]
show bgp [**ipv6** {**unicast**}] **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**]

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.
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 value4294967041 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

EXECEXEC mode

Command History

Release	Modification
Release 3.7.2	This command was introduced.
Release 5.3.2	The **graceful-shutdown** keyword was added.

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.

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 *System Management Command Reference for Cisco ASR 9000 Series Routers* 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 486 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 EXECEXEC mode mode:

```
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 11: 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.
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.

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.
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.

Related Commands

Command	Description
aggregate-address, on page 30	Creates an aggregate entry in a BGP routing table.
network (BGP), on page 236	Specifies a local network that the BGP routing process should originate and advertise to its neighbors.
route-policy (BGP), on page 290	Applies a routing policy to updates advertised to or received from a BGP neighbor
set default-afi	Sets the default Address Family Identifier (AFI) for the current session.
set default-safi	Sets the default Subaddress Family Identifier (SAFI) for the current session.
show bgp, on page 312	Displays BGP routes.

show bgp convergence

To display whether a specific address family has reached convergence, use the **show bgp convergence** command in EXEC mode.

show bgp [**ipv4** {**unicast** | **multicast** | **labeled-unicast** | **all** | **tunnel**}] **convergence**
show bgp [**ipv6** {**unicast**}] **convergence**
show bgp [**all** {**unicast** | **multicast** | **labeled-unicast** | **all** | **tunnel**}] **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.
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.

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

EXEC

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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.



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 *System Management Command Reference for Cisco ASR 9000 Series Routers* 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/RSP0/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/RSP0/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 12: 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.

Related Commands

Command	Description
set default-afi	Sets the default Address Family Identifier (AFI) for the current session.
set default-safi	Sets the default Subaddress Family Identifier (SAFI) for the current session.
show bgp summary, on page 496	Displays the status of BGP peer connections.
shutdown (BGP), on page 543	Disables a neighbor without removing all of its configuration.

show bgp dampened-paths

To display Border Gateway Protocol (BGP) dampened routes, use the **show bgp dampened-paths** command in EXECEXEC mode mode.

show bgp [**ipv4** {**unicast** | **multicast** | **labeled-unicast** | **all**}] **dampened-paths**
show bgp [**ipv6** {**unicast**}] **dampened-paths**
show bgp [**all** {**unicast** | **multicast** | **labeled-unicast** | **all** | **tunnel**}] **dampened-paths**
show bgp [**vpnv4 unicast** [**rd** *rd-address*]] **dampened-paths**
show bgp [**vrf** {*vrf-name* | **all**} [**ipv4** | {**unicast** | **labeled-unicast**} | **ipv6 unicast**]] **dampened-paths**

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

EXECEXEC mode

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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.



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 *System Management Command Reference for Cisco ASR 9000 Series Routers* 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 EXECEXEC mode mode:

```
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 13: 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.

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.
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.

Related Commands

Command	Description
aggregate-address, on page 30	Creates an aggregate entry in a BGP routing table.
bgp dampening, on page 86	Enables BGP route dampening or changes various BGP route dampening factors.
clear bgp dampening, on page 141	Clears BGP route dampening information and unsuppresses the suppressed routes.
network (BGP), on page 236	Specifies a local network that the BGP routing process should originate and advertise to its neighbors.
set default-afi	Sets the default Address Family Identifier (AFI) for the current session.
set default-safi	Sets the default Subaddress Family Identifier (SAFI) for the current session.
show bgp flap-statistics, on page 386	Displays BGP routes that have flapped.
show bgp neighbors, on page 417	Displays information about the TCP and BGP connections to neighbors.

show bgp flap-statistics

To display information about Border Gateway Protocol (BGP) paths that have flapped, use the **show bgp flap-statistics** command in EXECEXEC mode 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**}] **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**]

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

EXECEXEC mode

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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

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 *System Management Command Reference for Cisco ASR 9000 Series Routers* 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:

```
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 EXECEXEC mode mode:

```
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 14: 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.

Related Commands

Command	Description
bgp dampening, on page 86	Enables BGP route dampening or changes various BGP route dampening factors.
set default-afi	Sets the default Address Family Identifier (AFI) for the current session.
set default-safi	Sets the default Subaddress Family Identifier (SAFI) for the current session.
show bgp dampened-paths, on page 382	Displays the BGP dampened routes.
show bgp neighbors, on page 417	Displays information about BGP neighbors.

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 EXECEXEC mode mode.

show bgp [**ipv4** {**unicast** | **multicast** | **labeled-unicast** | **all** | **tunnel**}] **inconsistent-as** [**standby**]
show bgp [**ipv6** {**unicast**}] **inconsistent-as** [**standby**]
show bgp [**all** {**unicast** | **multicast** | **labeled-unicast** | **all** | **tunnel**}] **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**]

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

EXECEXEC mode

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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

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 *System Management Command Reference for Cisco ASR 9000 Series Routers* 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 EXECEXEC mode mode:

```
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 15: 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.

Related Commands

Command	Description
network (BGP), on page 236	Specifies a local network that the BGP routing process should originate and advertise to its neighbors.
route-policy (BGP), on page 290	Applies a routing policy to updates advertised to or received from a BGP neighbor.
set default-afi	Sets the default Address Family Identifier (AFI) for the current session.
set default -safi	Sets the default Subaddress Family Identifier (SAFI) for the current session.

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 EXEC mode.

show bgp l2vpn { **vpls** | **vpws** | **evpn** } { **recent-prefix** [**private**] | **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).
recent-prefixes	(Optional) Specifies summary of recently updated or deleted prefixes.
private	(Optional) Specifies private information.
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

EXEC mode

EXEC

Command History

Release	Modification
Release 3.9.1	This command was introduced.
Release 7.1.1	This command was modified. The keywords recent-prefixes and private were added.

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

The following example is sample output from the **show bgp l2vpn vpls** for route distinguisher 1:1:

```
RP/0/RSP0/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/RSP0/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/RSP0/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/RSP0/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/RSP0/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 recent-prefixes** command. This sample output provides details for the route distinguishers and corresponding prefix.

```
Router# show bgp l2vpn evpn  recent-prefixes

Address-Family   Route-Distinguisher Prefix
================ =================== ==============
L2VPN EVPN       0:0:0               [5][0][32][198.51.100.22]/24
L2VPN EVPN       10.5.0.1:100        [5][0][32][192.0.2.1]/24
L2VPN EVPN       10.5.0.1:100        [5][0][32][192.0.2.2]/24
L2VPN EVPN       10.5.0.1:100        [5][0][32][192.0.2.3]/24
L2VPN EVPN       10.5.0.1:100        [5][0][32][192.0.2.4]/24
```

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/RSP0/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 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

EXEC

Command History

Release	Modification
Release 3.9.1	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

Example

The following example is sample output from the **show bgp l2vpn vpls** command:

```
RP/0/RSP0/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/RSP0/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/RSP0/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/RSP0/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/RSP0/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/RSP0/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/RSP0/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 labels

To display Border Gateway Protocol (BGP) routes and their incoming and outgoing labels, use the **show bgp labels**command in 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

EXEC

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 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

Examples

The following is sample output from the **show bgp labels**command in EXEC mode:

```
RP/0/RSP0/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 16: 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.

Related Commands

Command	Description
set default-afi	Sets the default Address Family Identifier (AFI) for the current session.
set default-safi	Sets the default subaddress Family Identifier (SAFI) for the current session.

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 EXEC mode.

show bgp [**ipv4** { **unicast** | **multicast** | **labeled-unicast** | **all** | **tunnel** }] **large-community** *large-community-list* [**exact-match**]
show bgp [**ipv6** { **unicast** }] **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**]

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.
large-community	Specifies 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 large- communities exactly matching the specified large-communities.

Command Default

None.

Command Modes

EXEC

Command History

Release	Modification
Release 6.3.1	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.

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 EXEC mode:

```
RP/0/RSP0/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/RSP0/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.

Related Commands

Command	Description
route-policy (BGP), on page 290	Applies a routing policy to updates advertised to or received from a BGP neighbor.
show bgp, on page 312	Displays BGP routes.

show bgp memory history

To show the history of memory changes when memory reaches a threshold of more than 1% of rlimit (rlimit is the maximum memory usage of bgp before triggering bgp process self restart protection), use the **show bgp memory history** command in EXEC mode.

Note This command applies to all the Cisco IOS XR products running BGP.

show bgp memory history

Syntax Description

This command has no keywords or arguments.

Command Default

None

Command Modes

EXEC

Command History

Release	Modification
Release 7.11.1	This command was introduced.
Release 7.3.6	

Usage Guidelines

Use this command to monitor BGP memory history usage versus prefixes, paths, path elements (= best paths or advertised paths) and bgp attributes. If bgp memory usage keeps increasing over time while network items scale remains stable, this might be an indication of memory leak which needs further investigation.

Task ID

Task ID	Operations
bgp	read

The following example uses the **show bgp memory history** command to display syslogs when memory reaches the rlimit threshold.

```
**************************************************************************

RP/0/RP0/CPU0:Router#show bgp memory history

History of memory changes recorded for a threshold greater than 1.0% of rlimit.
Last shown record displays current values.
Network information for default VRF.

Time               Memory(MB)  Rlimit(%)  Memory diff(MB)  Networks  Paths     PathElems
Attributes
Oct  2 16:30:37    152         1          152              400       400       400
9
Oct  2 16:31:37    343         4          191              396952    396869    396952
725
Oct  2 16:32:37    425         5          81               524567    513979    524567
```

```
8408
...
Oct  3 21:39:49    818         9           277             1297568   1347868   1297555
11109
Oct  3 21:40:49    1081        13          263             1974078   2414610   1974065
11111
Oct  4 00:19:10    1081        13          -1              1972907   2413813   1972894
11154
```

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 EXECEXEC mode 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

EXECEXEC mode

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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

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:

```
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:

```
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:

```
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 17: 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.

Related Commands

Command	Description
af-group, on page 28	Configures a BGP address family group.
session-group, on page 307	Creates a session group and enters session group configuration mode.

Command	Description
show bgp af-group, on page 335	Displays information about configuration for address family groups.
show bgp neighbors, on page 417	Displays information about BGP neighbors, including configuration inherited from neighbor groups, session groups, and address family groups.
show bgp session-group, on page 490	Displays information about the BGP configuration for session groups.
show running-config	Displays the contents of the currently running configuration or a subset of that configuration.
use, on page 592	Inherits configuration from a neighbor group, a session group, or an address family group.

show bgp neighbors

To display information about Border Gateway Protocol (BGP) connections to neighbors, use the **show bgp neighbors** command in EXEC mode.

show bgp neighbors [**performance-statistics** | **missing-eor**] [**standby**]
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**]]
show bgp neighbor *ip-address* [**advertised-count** | **advertised-routes** | **configuration** | **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.
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.

advertised-count	(Optional) Displays the total number of advertised routes.
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.
decoded-message-log	(Optional) Displays BGP message logs.
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

EXEC

Command History

Release	Modification
Release 7.9.1	This command displays the slow peer configuration state and slow peer detection or processing information.
Release 3.7.2	This command was introduced.
Release 3.9.0	Asplain format for 4-byte autonomous system numbers notation was supported. The input parameters and output were modified to display 4-byte autonomous system numbers and extended communities in either asplain or asdot notations.
Release 5.3.2	The command was modified to include graceful maintenance feature information.
Release 5.3.2	The command output was modified display TCP MSS information.

Release	Modification
Release 5.3.2	The decoded-message-log [in | out] option was added.

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 *System Management Command Reference for Cisco ASR 9000 Series Routers* 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:

```
******************************************************************************

RP/0/0/CPU0:R1#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 12.12.12.5
 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:

```
RP/0/RSP0/CPU0: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 1.1.1.2
 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 18: 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.

Field	Description
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
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.

Field	Description
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.
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

Field	Description
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.
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.

Field	Description
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.
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

Field	Description
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:

```
RP/0/RSP0/CPU0: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:

```
RP/0/RSP0/CPU0: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 19: 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.

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.

The following is sample output from the **show bgp neighbors** command with the **dampened-routes** keyword:

```
RP/0/RSP0/CPU0: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 20: 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:

```
RP/0/RSP0/CPU0: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 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:

```
RP/0/RSP0/CPU0: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 22: 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:

```
RP/0/RSP0/CPU0: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 23: 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.

This sample output from the **show bgp neighbor** command displays TCP MSS information for the specified neighbor:

```
RP/0/RSP0/CPU0: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:

```
RP/0/RSP0/CPU0: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 []
```

This sample output from the **show bgp neighbor** command with the **advertised-count** keyword displays the number of prefixes advertised to a specific BGP neighbor:

```
Router# show bgp neighbor 10.10.0.1 advertised-count

Mon Aug 18 11:11:00.891 UTC
No of prefixes Advertised: 5
```

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*.

Related Commands

Command	Description
clear bgp, on page 137	Resets a BGP connection or session.
network (BGP), on page 236	Specifies a local network that the BGP routing process should originate and advertise to its neighbors.
route-policy (BGP), on page 290	Applies a routing policy to updates advertised to or received from a BGP neighbor.
set default-afi	Sets the default Address Family Identifier (AFI) for the current session.
set default-safi	Sets the default Subaddress Family Identifier (SAFI) for the current session.
show bgp, on page 312	Displays entries in the BGP routing table.

Command	Description
show bgp dampened-paths, on page 382	Displays BGP dampened routes.
show bgp flap-statistics, on page 386	Displays BGP routes that have flapped.
show bgp neighbor-group, on page 413	Displays information about the BGP configuration for neighbor groups.
shutdown (BGP), on page 543	Disables a neighbor without removing all of its configuration.

show bgp neighbors nsr

To display Border Gateway Protocol (BGP) nonstop routing (NSR) information across neighbors, use the **show bgp neighbors nsr** command in EXEC mode.

show bgp [**ipv4** {**unicast** | **multicast** | **all**} | **ipv6** {**unicast** | **multicast** | **all**} | **vpnv4 unicast** | **vpnv6 unicast** | **vrf** {**all***vrf_name*}] **neighbors nsr** [**standby**]

Syntax Description

ipv4	(Optional) Specifies IP Version 4 address prefixes.
ipv6	(Optional) Specifies IP Version 6 address prefixes.
unicast	(Optional) Specifies unicast address prefixes.
multicast	(Optional) Specifies multicast address prefixes.
all	(Optional) For address family, specifies prefixes for all address families.
vpnv4 unicast	(Optional) Specifies VPNv4 unicast address families.
vpnv6 unicast	(Optional) Specifies VPNv6 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.
standby	(Optional) Displays information about the standby card.

Command Default

No default behavior or values.

Command Modes

EXEC

Command History

Release	Modification
Release 3.9.0	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

Examples

The following is sample output from the **show bgp neighbors nsr** command with the **standby** keyword:

```
RP/0/RSP0/CPU0:router# show bgp neighbors nsr standby

BGP neighbor is 2.2.2.2
  BGP state = Established, up for 5d04h
  NSR state = NSR Ready
  Outstanding Postits: 0


BGP neighbor is 10.0.101.5
  BGP state = Established, up for 05:19:00
  NSR state = NSR Ready
  Outstanding Postits: 0


BGP neighbor is 10.1.0.5
  BGP state = Established, up for 5d04h
  NSR state = NSR Ready
  Outstanding Postits: 0
```

This table describes the significant fields shown in the display.

Table 24: show bgp neighbors nsr Field Descriptions

Field	Description
BGP state	Displays BGP neighbor peering state.
NSR state	Displays BGP neighbor NSR state.
Outstanding Postits	Displays the postit counters of pending events.

Related Commands

Command	Description
nsr (BGP), on page 249	Activates the Border Gateway Protocol (BGP) nonstop routing (NSR).
show bgp summary nsr, on page 500	Displays the Border Gateway Protocol (BGP) nonstop routing (NSR) information.
show bgp summary, on page 496	Displays the status of all Border Gateway Protocol (BGP) connections.

show bgp nexthops

To display statistical information about the Border Gateway Protocol (BGP) next hops, use the **show bgp nexthops** command in EXEC mode.

show bgp nexthops A.B.C.D.aigp-value[**statistics**] [**speaker** *speaker-id*] []

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

EXEC

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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

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. The 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/RSP0/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 25: 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.

Field	Description
Bestpaths Changed After Last Walk	Number of best paths modified due to the last notification.
Next Hop	IP address of the next hop.
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.

Related Commands

Command	Description
bgp redistribute-internal, on page 108	Specifies the delay for triggering BGP next-hop calculations.

show bgp nsr

To display Border Gateway Protocol (BGP) nonstop routing (NSR) information, use the **show bgp nsr** command in EXEC mode.

show bgp [**ipv4** {**unicast** | **multicast** | **labeled-unicast** | **all** | **tunnel**} | **ipv6** {**unicast** | **multicast** | **all** | **labeled-unicast**} | **all** {**unicast** | **multicast** | **all** | **labeled-unicast** | **tunnel**} | **vpnv4 unicast** | **vrf** {*vrf-name* | **all**} [**ipv4** {**unicast** | **labeled-unicast**} | **ipv6 unicast**] | **vpvn6 unicast**] **nsr** [**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.
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.
standby	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

EXEC

Command History

Release	Modification
Release 3.9.0	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

Examples

The following is sample output from the **show bgp nsr** command:

```
RP/0/RSP0/CPU0:router# show bgp nsr

Fri Jan 30 10:18:48.171 PST PDT

BGP Process Information:
BGP is operating in STANDALONE mode
Autonomous System: 100
Router ID: 10.1.0.1 (manually configured)
Default Cluster ID: 10.1.0.1
Active Cluster IDs:  10.1.0.1
Fast external fallover enabled
Neighbor logging is not enabled
Enforce first AS enabled
AS Path ignore is enabled
AS Path multipath-relax is enabled
Default local preference: 100
Default keepalive: 60
Graceful restart enabled
Restart time: 180
Stale path timeout time: 360
RIB purge timeout time: 600
Non-stop routing is enabled
Update delay: 120
Generic scan interval: 60

Address family: IPv4 Unicast
Dampening is not enabled
Client reflection is enabled in global config
Scan interval: 60
Main Table Version: 7034
IGP notification: IGPs notified
RIB has converged: version 1

========== Post Failover Summary for Active instance ==========

Node               Process            Read      Write    Inbound


node0_0_CPU0       Speaker          146.75      18.90       3.46

  Entered mode  Standby Ready              : Jan 30 10:00:39
  Entered mode  TCP NSR Setup              : Jan 30 10:00:39
```

```
  Entered mode  TCP NSR Setup Done          : Jan 30 10:00:39
  Entered mode  TCP Initial Sync            : Jan 30 10:00:39
  Entered mode  TCP Initial Sync Done       : Jan 30 10:00:44
  Entered mode  FPBSN processing done       : Jan 30 10:00:44
  Entered mode  Update processing done      : Jan 30 10:00:44
  Entered mode  BGP Initial Sync            : Jan 30 10:00:44
  Entered mode  BGP Initial Sync done       : Jan 30 10:00:44
  Entered mode  NSR Ready                   : Jan 30 10:00:44

Current BGP NSR state - NSR Ready achieved at: Jan 30 10:00:44
NSR State READY notified to Redcon at: Jan 30 10:16:58

NSR Post Failover Summary:


QAD Statistics:

  Messages Sent       : 512           ACKs Received       : 512
  Messages Received   : 8             ACKs Sent           : 8
  Send Failures       : 1             Send ACK Failures   : 0
  Suspends            : 1             Resumes             : 1
  Messages Processed  : 8             Out of sequence drops: 0

Postit Summary:

  Total pending postit messages:  0
  Neighbors with pending postits: 0



Conv  Bestpath   TunnelUpd  Import     RIBUpd     Label      ReadWrite  LastUpd
Process: Speaker

Yes   120        ---        ---        120        120        120        87531

Rib Trigger: enabled
Last RIB down event  Jan 29 09:50:03.069 received
Last RIB convergence Jan 29 09:50:03.069 last ack received.

Address Family IPv4 Unicast converged in 87531 seconds
```

The following example shows sample output from the **show bgp nsr** command with the **standby** keyword:

```
RP/0/RSP0/CPU0:router# show bgp nsr standby

Fri Jan 30 10:18:55.654 PST PDT

BGP Process Information:
BGP is operating in STANDALONE mode
Autonomous System: 100
Router ID: 10.1.0.1 (manually configured)
Default Cluster ID: 10.1.0.1
Active Cluster IDs:  10.1.0.1
Fast external fallover enabled
Neighbor logging is not enabled
Enforce first AS enabled
AS Path ignore is enabled
AS Path multipath-relax is enabled
Default local preference: 100
Default keepalive: 60
Graceful restart enabled
```

```
Restart time: 180
Stale path timeout time: 360
RIB purge timeout time: 600
Non-stop routing is enabled
Update delay: 120
Generic scan interval: 60

Address family: IPv4 Unicast
Dampening is not enabled
Client reflection is enabled in global config
Scan interval: 60
Main Table Version: 7034
IGP notification: IGPs notified
RIB has converged: version 1

========== Post Failover Summary for Standby instance ==========

Node                Process            Read      Write    Inbound


node0_1_CPU0        Speaker            1.68       0.00       1.42

Entered mode  Standby Ready              : Jan 30 10:00:39
Entered mode  TCP Replication            : Jan 30 10:00:39
Entered mode  TCP Init Sync Done         : Jan 30 10:00:44
Entered mode  NSR Ready                  : Jan 30 10:00:44

QAD Statistics:

  Messages Sent       : 9              ACKs Received       : 9
  Messages Received   : 512            ACKs Sent           : 512
  Send Failures       : 0              Send ACK Failures   : 0
  Suspends            : 0              Resumes             : 0
  Messages Processed  : 512            Standby init drops  : 0         Out of sequence
drops: 0

Postit Summary:

  Total pending postit messages:  0
  Neighbors with pending postits: 0



Conv  Bestpath   TunnelUpd  Import     RIBUpd     Label      ReadWrite  LastUpd
Process: Speaker

Yes   1233338444 ---        ---        1233338444 1233338444 1233338444 ---

Rib Trigger: enabled
Last RIB down event  Jan 29 09:50:17.308 received
Last RIB convergence Jan 29 09:50:17.308 last ack received.
```

Related Commands

Command	Description
nsr (BGP), on page 249	Activates Border Gateway Protocol (BGP) nonstop routing (NSR).

show bgp paths

To display all the Border Gateway Protocol (BGP) paths in the database, use the **show bgp paths** command in 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

EXEC

Command History

Release	Modification
Release 3.7.2	This command was introduced.
Release 3.9.0	Asplain format for 4-byte Autonomous system numbers notation was supported. The input parameters and output were modified to display 4-byte autonomous system numbers and extended communities in either asplain or asdot notations.

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.

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 *Cisco ASR 9000 Series Aggregation Services Router Getting Started Guide* 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/RSP0/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 26: 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.

Field	Description
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 EXECEXEC mode 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 EXECEXEC mode

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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

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 *System Management Command Reference for Cisco ASR 9000 Series Routers*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 EXECEXEC mode mode:

```
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 27: 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 EXECEXEC mode mode:

```
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 28: 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.

Related Commands

Command	Description
route-policy (BGP), on page 290	Applies an inbound or outbound routing policy to a neighbor.
set default-afi	Sets the default Address Family Identifier (AFI) for the current session.

Command	Description
set default-safi	Sets the default Subaddress Family Identifier (SAFI) for the current session.
show bgp advertised, on page 327	Displays routes advertised to neighbors.
show bgp neighbors, on page 417	Displays information about the TCP and BGP connections to neighbors.
show bgp route-policy, on page 486	Displays BGP information about networks that match an outbound route policy.

show bgp process

To display Border Gateway Protocol (BGP) process information, use the **show bgp process** command in 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**]

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

EXEC

Command History

Release	Modification
Release 3.7.2	This command was introduced.
Release 3.9.0	Asplain format for 4-byte autonomous system numbers notation was supported. The input parameters and output were modified to display 4-byte autonomous system numbers and extended communities in either asplain or asdot notations.
Release 5.3.2	The command output was modified to include graceful maintenance feature information.
Release 7.5.3	Updated output to display the delay of the BGP start-up process since the last router update.

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.



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 *System Management Command Reference for Cisco ASR 9000 Series Routers* 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/RSP0/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 29: 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.

Field	Description
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. 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.

Field	Description
Speaker	Speaker process. A speaker process is responsible for receiving, processing, and sending BGP messages to configured neighbors.
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/RSP0/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 30: 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.

Field	Description
Enforce first AS enabled	Indicates that strict checking of the first autonomous system (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 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.

Field	Description
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.
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/RSP0/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 31: 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.
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.

Field	Description
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.
Config	Real time (in seconds) spent processing configuration commands by this process.
Data	Real time (in seconds) spent providing operational data by this process.

Field	Description
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.

The following is sample output from the **show bgp process** command with the **performance-statistics** and **detail** keywords:

```
RP/0/RSP0/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 32: 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.
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.

Field	Description
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/RSP0/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 **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—
```

Related Commands

Command	Description
bgp bestpath as-path ignore, on page 62	Sets the autonomous system path length to ignore when calculating preferred paths.
bgp bestpath compare-routerid, on page 64	Compare identical routes received from external BGP (eBGP) peers during the best-path selection process and select the route with the lowest router ID.
bgp bestpath med always, on page 67	Compare the Multi Exit Discriminator (MED) for paths from neighbors in different autonomous systems.
bgp bestpath med missing-as-worst, on page 71	Assume paths with no MED attribute have the most undesirable MED value possible when performing path selection.
bgp cluster-id, on page 80	Enables reflection of routes between route reflector clients using a BGP route reflector.
bgp cluster-id, on page 80	Configure the cluster ID if the BGP cluster has more than one route reflector.
bgp default local-preference, on page 88	Sets the default local preference value.
bgp redistribute-internal, on page 108	Allows the redistribution of iBGP routes into an IGP such as Intermediate System-to-Intermediate System (IS-IS) or Open Shortest Path First (OSPF).
bgp router-id, on page 110	Configures a fixed router ID for a BGP-speaking router.
default-metric (BGP), on page 166	Sets default metric values for the BGP.
set default-afi	Sets the default Address Family Identifier (AFI) for the current session.
set default-safi	Sets the default Subaddress Family Identifier (SAFI) for the current session.
bgp scan-time, on page 111	Configures scanning intervals.
timers bgp, on page 569	Sets default BGP timers.

show bgp regexp

To display routes matching the autonomous system path regular expression, use the**show bgp regexp** command in EXECEXEC mode 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.
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.
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

EXECEXEC mode

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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

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 *System Management Command Reference for Cisco ASR 9000 Series Routers* 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:

```
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 33: 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.

Related Commands

Command	Description
set default-afi	Sets the default Address Family Identifier (AFI) for the current session.
set default-safi	Sets the default Subaddress Family Identifier (SAFI) for the current session.
show bgp, on page 312	Displays entries in the BGP routing table.
show bgp route-policy, on page 486	Displays BGP information about networks that match an outbound route policy.

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 EXECEXEC mode mode.

show bgp route-policy *route-policy-name* []

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

EXECEXEC mode

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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.



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 *System Management Command Reference for Cisco ASR 9000 Series Routers* 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 EXECEXEC mode mode:

```
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 34: show bgp route-policy 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.

Related Commands

Command	Description
aggregate-address, on page 30	Configures an aggregate entry in a BGP routing table.
network (BGP), on page 236	Specifies a local network that the BGP routing process should originate and advertise to its neighbors.
route-policy (BGP), on page 290	Applies a routing policy to updates advertised to or received from a BGP neighbor
route-policy	Configures a route policy.
set default-afi	Sets the default Address Family Identifier (AFI) for the current session.
set default-safi	Sets the default Subaddress Family Identifier (SAFI) for the current session.
show bgp policy, on page 454	Displays advertisements under a proposed policy.

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 EXECEXEC mode 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

EXECEXEC mode

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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

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:

```
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:

```
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.

```
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 35: 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.
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.

Related Commands

Command	Description
session-group, on page 307	Configures a BGP session group.
show bgp neighbor-group, on page 413	Displays information about the BGP configuration for neighbor groups.
show bgp neighbors, on page 417	Displays information about BGP connections to neighbors.

show bgp sessions

To display brief information about BGP neighbors, use the **show bgp sessions** command in EXEC mode.

show bgp sessions [**not-established**] [**not-nsr-ready**]

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.

Command Default

No default behavior or values

Command Modes

EXEC

Command History

Release	Modification
Release 3.9.0	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.

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 EXEC mode:

```
RP/0/RSP0/CPU0:router# show bgp sessions
Thu Jan 15 17:41:45.277 UTC

Neighbor        VRF                  Spk     AS    InQ  OutQ  NBRState     NSRState
2.2.2.2         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
90.0.0.2        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
20.0.101.1      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
1.2.3.4         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:

```
RP/0/RSP0/CPU0: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
2.2.2.2         vrf1_1                  0   302     0     0  Idle         None
2.101.1.2       vrf1_1                  0   302     0     0  Idle         None
2.102.1.2       vrf1_1                  0   302     0     0  Idle         None
2.103.1.2       vrf1_1                  0   302     0     0  Idle         None
4.4.4.2         vrf1_1                  0   304     0     0  Idle         None
2008:2:2:2::2   vrf1_1                  0   302     0     0  Idle         None
11.16.1.2       vrf2_1                  0   302     0     0  Idle         None
```

The following is sample output from the **show bgp sessions** command with the **not-nsr-ready** keyword:

```
RP/0/RSP0/CPU0: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
2.2.2.2         vrf1_1                  0   302     0     0  Idle         None
2.101.1.2       vrf1_1                  0   302     0     0  Idle         None
2.102.1.2       vrf1_1                  0   302     0     0  Idle         None
2.103.1.2       vrf1_1                  0   302     0     0  Idle         None
```

```
4.4.4.2         vrf1_1                  0   304     0     0  Idle          None
2008:2:2:2::2   vrf1_1                  0   302     0     0  Idle          None
11.16.1.2       vrf2_1                  0   302     0     0  Idle          None
```

This table describes the significant fields shown in the display.

Table 36: 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).

Related Commands

Command	Description
show bgp neighbors, on page 417	Displays information about Border Gateway Protocol (BGP) connections to neighbors.

show bgp summary

To display the status of all Border Gateway Protocol (BGP) connections, use the **show bgp summary** command in EXECEXEC mode 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

EXECEXEC mode

Command History

Release	Modification
Release 3.7.2	This command was introduced.

Release	Modification
Release 3.9.0	Asplain format for 4-byte Autonomous system numbers notation was supported. The input parameters and output were modified to display 4-byte autonomous system numbers and extended communities in either asplain or asdot notations.

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.



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 *System Management Command Reference for Cisco ASR 9000 Series Routers* 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:

```
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 37: show bgp summary 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 295 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.

Field	Description
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.
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.

Related Commands

Command	Description
route-policy (BGP), on page 290	Applies a routing policy to updates advertised to or received from a BGP neighbor.
set default-afi	Sets the default Address Family Identifier (AFI) for the current session.
set default-safi	Sets the default Subaddress Family Identifier (SAFI) for the current session.

show bgp summary nsr

To display the summary of Border Gateway Protocol (BGP) neighbor state and nonstop routing (NSR) state information, use the **show bgp summary nsr** command in EXEC mode.

show bgp summary [**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**] **nsr** [**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.
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.
standby	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

EXEC

Command History

Release	Modification
Release 3.9.0	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

Examples

The following is sample output from the **show bgp summary nsr** command:

```
RP/0/RSP0/CPU0:router# show bgp summary nsr

BGP router identifier 10.1.0.1, local AS number 100
BGP generic scan interval 60 secs
Non-stop routing is enabled
BGP table state: Active
Table ID: 0xe0000000
BGP main routing table version 13037
BGP NSR Initial initsync version 11034 (Reached)
BGP scan interval 60 secs

BGP is operating in STANDALONE mode.


node0_1_CPU0        Speaker

Entered mode  Standby Ready              : Feb  3 14:22:00
Entered mode  TCP NSR Setup              : Feb  3 14:22:00
Entered mode  TCP NSR Setup Done         : Feb  3 14:22:01
Entered mode  TCP Initial Sync           : Feb  3 14:22:01
Entered mode  TCP Initial Sync Done      : Feb  3 14:22:44
Entered mode  FPBSN processing done      : Feb  3 14:22:44
Entered mode  Update processing done     : Feb  3 14:22:44
Entered mode  BGP Initial Sync           : Feb  3 14:22:44
Entered mode  BGP Initial Sync done      : Feb  3 14:22:49
Entered mode  NSR Ready                  : Feb  3 14:22:49

Current BGP NSR state - NSR Ready achieved at: Feb  3 14:22:49
NSR State READY notified to Redcon at: Feb  4 07:44:43

Process       RcvTblVer   bRIB/RIB   LabelVer  ImportVer  SendTblVer  StandbyVer
Speaker           13037      13037      13037      13037       13037       13037

Neighbor        Spk    AS   TblVer  SyncVer   AckVer NBRState     NSRState
2.2.2.2           0   302    13037    13037    13037 Established  NSR Ready
10.0.101.5        0   100    13037    13037    13037 Established  NSR Ready
```

The following example shows sample output from the **show bgp summary nsr** command with the **standby** keyword:

```
RP/0/RSP0/CPU0:router# show bgp summary nsr standby
```

```
BGP router identifier 10.1.0.1, local AS number 100
BGP generic scan interval 60 secs
Non-stop routing is enabled
BGP table state: Active
Table ID: 0xe0000000
BGP main routing table version 13037
BGP NSR Initial initsync version 0 (Not Reached)
BGP scan interval 60 secs

BGP is operating in STANDALONE mode.


node0_0_CPU0        Speaker

Entered mode  Standby Ready               : Feb  3 14:22:03
Entered mode  TCP Replication             : Feb  3 14:22:03
Entered mode  TCP Init Sync Done          : Feb  3 14:22:47
Entered mode  NSR Ready                   : Feb  3 14:22:52

Process       RcvTblVer   bRIB/RIB   LabelVer  ImportVer  SendTblVer  StandbyVer
Speaker           13037          0          0      13037           0           0

Neighbor        Spk    AS   TblVer  SyncVer   AckVer NBRState     NSRState
2.2.2.2           0   302    13037        0        1 Established  NSR Ready
10.0.101.5        0   100    13037        0        1 Established  NSR Ready
```

This table describes the significant fields shown in the display.

Table 38: show bgp summary nsr Field Descriptions

Field	Description
BGP router identifier	IP address of the router.
BGP generic scan interval	Interval (in seconds) between scans of the BGP table by a generic scanner.
Non-stop routing	State of the Nonstop routing.
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.
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.
Entered mode	The successive transition of various states of TCP and BGP, leading to the NSR ready state. **Note** This is used for monitoring and debugging purposes.
SyncVer	The version which has synced to standby for this neighbor.

Field	Description
AckVer	The version which the neighbor has acknowledge.
NBRState	State of the BGP neighbor.
NSRState	Neighbor NSR state.

Related Commands

Command	Description
nsr (BGP), on page 249	Activates Border Gateway Protocol (BGP) nonstop routing (NSR)
show bgp nsr, on page 447	Displays Border Gateway Protocol (BGP) nonstop routing (NSR) information.

show bgp table

To display the status of all Border Gateway Protocol (BGP) neighbors for a particular Address Family (AF) in the global address table, use the **show bgp table** command in EXEC mode.

show bgp table [**ipv4** {**mdt** | **multicast** | **mvpn** | **rt-filter** | **tunnel** | **unicast**} | **ipv6** {**multicast** | **mvpn** | **unicast**} | **l2vpn** {**evpn** | **vpls** | **vpws**} | **standby** | **vpnv4 unicast** | **vpnv6 unicast**]

Syntax Description

ipv4 mdt	(Optional) Specifies IPv4 multicast distribution tree (MDT) neighbors.
ipv4 multicast	(Optional) Specifies IPv4 multicast neighbors.
ipv4 mvpn	(Optional) Specifies the IPv4 mvpn address family neighbors.
ipv4 rt-filter	(Optional) Specifies the IPv4 RT Constraint address family neighbors.
ipv4 tunnel	(Optional) Specifies IPv4 tunnel neighbors.
ipv6 unicast	(Optional) Specifies IP Version 6 (IPv6) unicast neighbors.
ipv6 multicast	(Optional) Specifies IPv6 multicast neighbors.
ipv6 mvpn	(Optional) Specifies the IPv6 mvpn address family neighbors.
ipv6 unicast	(Optional) Specifies the IPv6 Tunnel address family neighbors.
l2vpn evpn	(Optional) Specifies the L2VPN EVPN address family neighbors.
l2vpn vpls	(Optional) Specifies the L2VPN VPLS address family neighbors.
l2vpn vpws	(Optional) Specifies the L2VPN VPWS address family neighbors.
standby	(Optional) Specifies the IPv4 Unicast address family neighbor on the standby processor.
vpnv4 unicast	(Optional) Specifies VPN Version 4 (VPNv4) unicast address family neighbors.
vpnv6 unicast	(Optional) Specifies VPN Version 6 (VPNv6) unicast address family neighbors.

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

EXEC

Command History

Release	Modification
Release 3.8.0	This command was introduced.
Release 4.3.2, 5.1.0, 5.1.1, 5.1.2 and 5.2.0	The L2VPN Address Family was added.

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.

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 *Cisco ASR 9000 Series Aggregation Services Router System Management Command Reference* 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 table** command to display a brief summary of the neighbors for which the specified address family (AFI) and subaddress family (SAFI) are enabled. If the AFI and/or SAFI is not enabled, the command will only display the column headings.

Task ID

Task ID	Operations
bgp	read

Examples

The following is sample output from the **bgp table vpnv4 unicast** command in EXEC mode:

```
RP/0/RSP0/CPU0:router# show bgp table vpnv4 unicast
Thu Jan 15 17:43:31.215 UTC
Neighbor        VRF                  Spk     AS   TblVer  InQ OutQ  St/PfxRcd
10.0.101.1      default                0      1      951    0    0          11
10.0.101.2      default                0      1      951    0    0           5
10.0.101.3      default                0      1      951    0    0           0
10.0.101.4      default                0      1      951    0    0           0
10.0.101.5      default                0      1      951    0    0           0
10.0.101.6      default                0      1      951    0    0           0
10.0.101.7      default                0      1      951    0    0           0
10.0.101.8      default                0      1      951    0    0           0
10.0.101.9      default                0      1      951    0    0           0
90.0.0.2        900                    0      2      951    0    0           1
91.0.0.2        901                    0      2      951    0    0           1
92.0.0.2        902                    0      2      951    0    0           1
93.0.0.2        903                    0      2      951    0    0           3
94.0.0.2        904                    0      2      951    0    0           3
95.0.0.2        905                    0      2      951    0    0           3
96.0.0.2        906                    0      2      951    0    0           3
97.0.0.2        907                    0      2      951    0    0           3
98.0.0.2        908                    0      2      951    0    0           3
99.0.0.2        909                    0      2        0    0    0 Idle
12.13.14.16     red                    0      2        0    0    0 Idle
20.0.101.1      red                    0      2        0    0    0 Active
1.2.3.4         this-is-a-long-vrf-name
                                       0      5        0    0    0 Idle
```

This table describes the significant fields shown in the display.

Table 39: show bgp table Field Descriptions

Field	Description
Neighbor	IP address of a neighbor.

Field	Description
VRF	The VRF which each neighbor belongs to; either the default VRF or a specified VRF.
Spk	Speaker process that is responsible for the neighbor. Always 0.
AS	Autonomous system.
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.
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.

Related Commands

Command	Description
show bgp neighbor-group, on page 413	Displays information about the Border Gateway Protocol (BGP) configuration for neighbor groups.
show bgp neighbors, on page 417	Displays information about Border Gateway Protocol (BGP) connections to neighbors.
show bgp summary, on page 496	Displays the status of all Border Gateway Protocol (BGP) connections.

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 EXEC mode mode.

show bgptruncated-communities

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.
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

EXEC mode

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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

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 *System Management Command Reference for Cisco ASR 9000 Series Routers* 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:

```
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 40: 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.

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.

Related Commands

Command	Description
aggregate-address, on page 30	Creates an aggregate entry in a BGP routing table.
network (BGP), on page 236	Specifies a local network that the BGP routing process should originate and advertise to its neighbors.
route-policy (BGP), on page 290	Applies a routing policy to updates advertised to or received from a BGP neighbor.
set default-afi	Sets the default Address Family Identifier (AFI) for the current session.
set default-safi	Sets the default Subaddress Family Identifier (SAFI) for the current session.
show bgp, on page 312	Displays entries in the BGP routing table.

show bgp update-group

To display Border Gateway Protocol (BGP) information for update groups, use the **show bgp update-group** command in 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

EXEC EXEC

Command History

Release	Modification
Release 3.7.2	This command was introduced.
Release 3.9.0	Asplain format for 4-byte Autonomous system numbers notation was supported. The input parameters and output were modified to display 4-byte autonomous system numbers and extended communities in either asplain or asdot notations.
Release 7.1.1	The command output was modified to include the performance statistics for BGP peer that was not advertised..

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 *System Management Command Reference for Cisco ASR 9000 Series Routers* 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 496 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/RSP0/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
```

The following is sample output from the **show bgp update-group performance-statistics** command:

```
RP/0/RSP0/CPU0:router# show bgp update-group performace-statistics


Update group for IPv4 Unicast, index 0.1:
..
Update timer last processed: Sep 23 00:10:15.350
  Not-Advertised Stats:
    Non-Client to Non-Client     : 105        Sep 23 00:10:15.350
    Path Not Available     : 132        Sep 23 00:10:15.350
```

The following is sample output from the **show bgp update-group performance-statistics** command:

```
Router# show bgp update-group performace-statistics


Update group for IPv4 Unicast, index 0.1:
..
Update timer last processed: Sep 23 00:10:15.350
  Not-Advertised Stats:
    Non-Client to Non-Client     : 105        Sep 23 00:10:15.350
    Path Not Available     : 132        Sep 23 00:10:15.350
```

This table describes the significant fields shown in the display.

Table 41: 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.

Field	Description
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.
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/RSP0/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 table describes the significant fields shown in the display.

Table 42: 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 295 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.

Field	Description
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.
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.

This is sample output from the **show bgp update-group performance-statistics** command that shows the total numbers of prefixes not advertising in any AFI or SAFI, including repeating counts on 1 or more prefixes:

```
Router# show bgp update-group performance-statistics

Update group for IPv4 Unicast, index 0.1:
..
Update timer last processed: Sep 23 00:10:15.350
  Not-Advertised Stats:
    Non-Client to Non-Client       : 105        Sep 23 00:10:15.350
    Path Not Available       : 132        Sep 23 00:10:15.350
```

Related Commands

Command	Description
maximum-prefix (BGP), on page 221	Limits the number of prefixes that can be received from a neighbor.
route-policy (BGP), on page 290	Applies a routing policy to updates advertised to or received from a BGP neighbor.
set default-afi	Sets the default Address Family Identifier (AFI) for the current session.
set default-safi	Sets the default Subaddress Family Identifier (SAFI) for the current session.
show bgp summary, on page 496	Displays the status of all BGP connections.
shutdown (BGP), on page 543	Disables a neighbor without removing its configuration.

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 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

EXEC

Command History

Release	Modification
Release 4.2.0	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

This example displays sample output from the **show bgp update in error process** command:

```
RP/0/RSP0/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

To display address-family level update generation information, use the **show bgp update out** command in 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

EXEC

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 4.2.0	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

This example displays sample output from the **show bgp update out** command:

```
RP/0/RSP0/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 10.1.114.1 detail

Tue Nov 15 18:48:24.863 UTC

VRF "default", Address-family "VPNv6 Unicast"
  Main routing table version: 51001
  RIB version: 51001
```

```
  Neighbor 10.1.114.1
    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 out filter-group

To display update generation information at filter-group level, **show bgp update out filter-group** command in 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

EXEC

Command History

Release	Modification
Release 4.2.0	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

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 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

EXEC

Command History

Release	Modification
Release 4.2.0	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

This example displays sample output from the show bgp update out process brief command:

```
RP/0/RSP0/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 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

EXEC

Command History

Release	Modification
Release 4.2.0	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

This example displays sample output from the **show bgp update out sub-group** command:

```
RP/0/RSP0/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 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

EXEC

Command History

Release	Modification
Release 4.2.0	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

This example shows the significant fields on display form the **show bgp update out update-group** command:

```
RP/0/RSP0/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

To display Border Gateway Protocol (BGP) prefix information for VPN routing and forwarding (VRF) instances, use the **show bgp vrf** command in EXEC mode.

show bgp vrf { **all***vrf-name* } { **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 routes.
ipv4-address/length [**detail**]	(Optional) Displays detailed output for the specified IPv4-address.
ipv6 unicast	(Optional) Specifies IP Version 6 unicast 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.

Command Default

No default behavior or values

Command Modes

EXEC mode

Command History

Release	Modification
Release 3.7.2	This command was introduced.
Release 3.9.0	Asplain format for 4-byte Autonomous system numbers notation was supported. The input parameters and output were modified to display 4-byte autonomous system numbers and extended communities in either asplain or asdot notations.

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.

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/RSP0/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
```

This table describes the significant fields shown in the display output for **show bgp vrf** command.

Table 43: 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 295 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.

Field	Description
Route Distinguisher	Unique identifier for the routing instance.
Source VRF	Source VRF for the imported route.

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 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

EXEC

Command History

Release	Modification
Release 4.2.0	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

This example displays sample output from s**how bgp vrf vrf1 update in error** command:

```
RP/0/RSP0/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 lpts punt excessive-flow-trap bgp

To display the details of bad actor identified for bgp protocol, use the **show lpts punt excessive-flow-trap bgp** command in the Global Configuration mode.

show lpts punt excessive-flow-trap bgp

Command Default

None

Command Modes

Global Configuration mode

Command History

Release	Modification
Release 6.0.1	This command was introduced.

Usage Guidelines

You must be in a user group associated with a task group that includes the proper task IDs. The command reference guides include the task IDs required for each command. If you suspect user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Task ID

Task ID	Operations
lpts	read
basic-services	read-write

Examples

This example shows to the details of bad actor identified for bgp protocol:

```
RP/0/RSP0/CPU0:router# show lpts punt excessive-flow-trap bgp
Parent Interface: GigabitEthernet0/2/0/19          Src MAC Addr: 0000.6505.0102

         Intf Handle: 0x08000580                          Location: 0/2/CPU0
            Protocol: BGP                                      Punt Reason: BGP-mc-known

        Penalty Rate: 0 pps (all packets dropped)    Penalty Timeout: 15 mins

      Time Remaining: 10 mins 3 secs
```

Related Commands

Command	Description
show running-config lpts punt excessive-flow-trap, on page 537	Displays the running configuration for the Excessive Punt Flow Trap feature.

show protocols (BGP)

To display information about the Border Gateway Protocol (BGP) instances running on the router, use the **show protocols** command in 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** , **eigrp** , and **ospf** . For the IPv6 address family, the options are **bgp** , **eigrp** , **isis** , and **ospfv3** .

Command Default

Default is IPv4.

Command Modes

EXEC

Command History

Release	Modification
Release 3.7.2	This command was introduced.
Release 3.9.0	Asplain format for 4-byte Autonomous system numbers notation was supported. The input parameters and output were modified to display 4-byte autonomous system numbers and extended communities in either asplain or asdot notations.

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.

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/RSP0/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 44: 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 running-config lpts punt excessive-flow-trap

To display the running configuration for the Excessive Punt Flow Trap feature, use the **show running-config lpts punt excessive-flow-trap** command in the EXEC mode.

show running-configlpts punt excessive-flow-trap

Command Default

None

Command Modes

EXEC mode

Command History

Release	Modification
Release 5.2.2	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

Task ID

Task ID	Operation
lpts	read
basic-services	read, write

The **show running-config** output for the above **show lpts punt excessive-flow-trap** command is:

```
RP/0/RSP0/CPU0:router# show running-config lpts punt excessive-flow-trap
lpts punt excessive-flow-trap
penalty-rate arp 10
penalty-timeout arp 20
non-subscriber-interfaces
!
```

This table describes the significant fields shown in the display.

Table 45: show lpts punt excessive-flow-trap Field Descriptions

Field	Description
penalty-rate	The penalty policing rate for the ARP protocol. For ARP, the value is 10.
penalty-timeout	The penalty timeout value for the ARP protocol. For ARP, the value is 20.

Related Commands

Command	Description
lpts punt excessive-flow-trap	
show lpts punt excessive-flow-trap information	

Command	Description
show lpts punt excessive-flow-trap interface	

show svd role

To display selective VRF download (SVD) role information, use the **show svd role** command in EXEC mode.

show svd role

Syntax Description

This command has no keywords or arguments.

Command Default

None.

Command Modes

EXEC

Command History

Release	Modification
Release 4.2.0	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.

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/RSP0/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
```

Related Commands

Command	Description
selective-vrf-download disable, on page 299	Disables selective VRF Download (SVD).
show svd state, on page 540	Displays Selective VRF Download (SVD) state information.

show svd state

To display selective VRF download (SVD) state information, use the **show svd state** command in EXEC mode.

show svd state

Syntax Description

This command has no keywords or arguments.

Command Default

None.

Command Modes

EXEC

Command History

Release	Modification
Release 4.2.0	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
ip-services	read

This example shows the SVD configuration state and the SVD operational state in a line card:

```
RP/0/RSP0/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
```

Related Commands

Command	Description
selective-vrf-download disable, on page 299	Disables selective VRF Download (SVD).
show svd role, on page 539	Displays Selective VRF Download (SVD) role information.

show tcp pcb

To display TCP connection information, use the **show tcp pcb** command in the EXEC mode.

show tcp pcb *pcb-value*

Syntax Description

pcb-value	Specifies PCB hexadecimal address. The valid range is from 0x0 to 0xFFFFFFFF.

Command Default

No default behavior or values

Command Modes

EXEC mode

Command History

Release	Modification
Release 5.3.2	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 is a sample output from the **show tcp pcb** command:

```
RP/0/0/CPU0:ios#show tcp pcb 0xecb0c9f8

Connection state is ESTAB, I/O status: 0, socket status: 0
Established at Sun Dec 7 11:49:39 2014

PCB 0xecb0c9f8, SO 0xecb01b68, TCPCB 0xecb01d78, vrfid 0x60000000,
Pak Prio: Medium, TOS: 192, TTL: 255, Hash index: 1322
Local host: 10.0.0.1, Local port: 12404 (Local App PID: 19840)
Foreign host: 10.0.0.2, Foreign port: 179

Current send queue size in bytes: 0 (max 24576)
Current receive queue size in bytes: 0 (max 32768) mis-ordered: 0 bytes
Current receive queue size in packets: 0 (max 0)

Timer Starts Wakeups Next(msec)
Retrans 17 2 0
SendWnd 0 0 0
TimeWait 0 0 0
AckHold 13 5 0
KeepAlive 1 0 0
PmtuAger 0 0 0
GiveUp 0 0 0
Throttle 0 0 0
```

```
iss: 1728179225 snduna: 1728179536 sndnxt: 1728179536
sndmax: 1728179536 sndwnd: 32517 sndcwnd: 1000
irs: 2055835995 rcvnxt: 2055836306 rcvwnd: 32536 rcvadv: 2055868842

SRTT: 206 ms, RTTO: 300 ms, RTV: 59 ms, KRTT: 0 ms
minRTT: 10 ms, maxRTT: 230 ms

ACK hold time: 200 ms, Keepalive time: 0 sec, SYN waittime: 30 sec
Giveup time: 0 ms, Retransmission retries: 0, Retransmit forever: FALSE
Connect retries remaining: 30, connect retry interval: 30 secs

State flags: none
Feature flags: Win Scale, Nagle
Request flags: Win Scale

Datagrams (in bytes): MSS 500, peer MSS 1460, min MSS 500, max MSS 1460

Window scales: rcv 0, snd 0, request rcv 0, request snd 0
Timestamp option: recent 0, recent age 0, last ACK sent 0
Sack blocks {start, end}: none
Sack holes {start, end, dups, rxmit}: none

Socket options: SO_REUSEADDR, SO_REUSEPORT, SO_NBIO
Socket states: SS_ISCONNECTED, SS_PRIV
Socket receive buffer states: SB_DEL_WAKEUP
Socket send buffer states: SB_DEL_WAKEUP
Socket receive buffer: Low/High watermark 1/32768
Socket send buffer : Low/High watermark 2048/24576, Notify threshold 0

PDU information:
#PDU's in buffer: 0
FIB Lookup Cache: IFH: 0x200 PD ctx: size: 0 data:
Num Labels: 0 Label Stack:
```

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 3.7.2	This command was introduced.
Release 3.9.0	The **disable** keyword was replaced with the **inheritance-disable** keyword.

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.

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:

```
Router(config)# router bgp 1
Router(config-bgp)# neighbor 192.168.40.24
Router(config-bgp-nbr)# shutdown
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:

```
Router(config)# router bgp 1
Router(config-bgp)# session-group group1
Router(config-bgp-sngrp)# shutdown
Router(config-bgp-sngrp)# exit
Router(config-bgp)# neighbor 192.168.40.24
Router(config-bgp-nbr)# remote-as 1
Router(config-bgp-nbr)# use session-group group1
Router(config-bgp-nbr)# shutdown inheritance-disable
Router(config-bgp-nbr)# exit
```

Related Commands

Command	Description
neighbor-group, on page 231	Creates a neighbor group and enters neighbor group configuration mode.
session-group, on page 307	Creates a session group and enters session group configuration mode.
show bgp summary, on page 496	Displays the status of all BGP connections.

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 4.2.1	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

This command shows how to configure no shutdown of the RPKi cache configuration after other RPKI cache parameters are configured:

```
Router#configure
Router(config)#router bgp 100
Router(config-bgp)#rpki server 172.168.35.40
Router(config-bgp-rpki-server)# transport ssh port 22
Router(config-bgp-rpki-server)#username rpki-user
Router(config-bgp-rpki-server)#password rpki-ssh-pass
Router(config-bgp-rpki-server)#preference 1
Router(config-bgp-rpki-server)#purge-time 30
Router(config-bgp-rpki-server)#refresh-time 30
Router(config-bgp-rpki-server)#response-time 30
Router(config-bgp-rpki-server)#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 3.9.1	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 disable BGP signaling protocol for neighbor 10.2.3.4:

```
RP/0/RSP0/CPU0:router(config)#router bgp 100
RP/0/RSP0/CPU0:router(config-bgp)#neighbor 10.2.3.4
RP/0/RSP0/CPU0:router(config-bgp-nbr)#address-family l2vpn vpls-vpws
RP/0/RSP0/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 3.7.2	This command was introduced.
Release 3.9.0	Asplain format for 4-byte Autonomous system numbers notation was supported.

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.

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/RSP0/CPU0:router(config)# router bgp 6
RP/0/RSP0/CPU0:router(config-bgp)# vrf vrf_A
RP/0/RSP0/CPU0:router(config-bgp-vrf)# neighbor 192.168.70.24
RP/0/RSP0/CPU0:router(config-bgp-vrf-nbr)# remote-as 10
RP/0/RSP0/CPU0:router(config-bgp-vrf-nbr)# address-family ipv4 unicast
RP/0/RSP0/CPU0:router(config-bgp-vrf-nbr-af)# site-of-origin 10.0.01:20
```

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
```

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
```

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 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

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/RSP0/CPU0:router(config)# router bgp 1
RP/0/RSP0/CPU0:router(config-bgp)# socket receive-buffer-size 65536 8192
```

Related Commands

Command	Description
receive-buffer-size, on page 269	Sets the size of the receive buffers for a BGP neighbor.
socket send-buffer-size, on page 556	Sets the size of the send buffers for all BGP neighbors.

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 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

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/RSP0/CPU0:router(config)# router bgp 1
RP/0/RSP0/CPU0:router(config-bgp)# socket send-buffer-size 8192 8192
```

Related Commands

Command	Description
send-buffer-size, on page 300	Sets the size of the send buffers for a BGP neighbor.
socket receive-buffer-size, on page 554	Sets the size of the receive buffers for all BGP neighbors.

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

VRF IPv6 neighbor address family configuration

Command History

Release	Modification
Release 3.7.2	This command was introduced.
Release 3.9.0	The **disable** keyword was replaced with the **inheritance-disable** keyword.

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.

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.

```
Router(config)# router bgp 100
Router(config-bgp)# neighbor 10.108.1.1
Router(config-bgp-nbr)# remote-as 100
Router(config-bgp-nbr)# address-family ipv4 unicast
Router(config-bgp-nbr-af)# soft-reconfiguration inbound
Router(config-bgp-nbr-af)# exit
```

Related Commands

Command	Description
af-group, on page 28	Creates an address family group for BGP neighbors and enters address family group configuration mode.
bgp auto-policy-soft-reset disable, on page 60	Disables an automatic soft reset of BGP peers when the configured inbound route policy is modified.
clear bgp, on page 137	Resets a BGP connection using a soft or hard reset.
neighbor-group, on page 231	Creates a neighbor group and enters neighbor group configuration mode.
rd, on page 267	Applies a prefix list to filter updates received from a neighbor.
route-policy (BGP), on page 290	Applies a routing policy to updates advertised to or received from a BGP neighbor.

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 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 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

The following example shows how to allocate speaker process 3 to neighbor 192.168.40.24:

```
RP/0/RSP0/CPU0:router(config)# router bgp 109
RP/0/RSP0/CPU0:router(config-bgp)# neighbor 192.168.40.24
RP/0/RSP0/CPU0:router(config-bgp-nbr)# speaker-id 3
```

Related Commands

Command	Description
distributed speaker, on page 174	Starts a specified speaker process.

svd platform enable

To enable selective VRF download (SVD) for Cisco ASR 9000 Series Aggregation Services Router, use the **svd platform enable** command in administrative configuration 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

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.

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 Cisco ASR 9000 Series Aggregation Services Router, 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

Administrative configuration

Command History

Release	Modification
Release 4.3.1	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	Operation
config-services	read, write

This example shows how to enable selective VRF download for Cisco ASR 9000 Series Aggregation Services Router:

```
Router#configure
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 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

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/RSP0/CPU0:router(config)# router bgp 1
RP/0/RSP0/CPU0:router(config-bgp)# address-family ipv4 unicast
RP/0/RSP0/CPU0:router(config-bgp-af)# table-policy set-traffic-index
```

Related Commands

Command	Description
route-policy (RPL)	Defines a route policy and enters route policy configuration mode.

tcp mss inheritance-disable

To disable TCP MSS under neighbor or neighbor group, or to prevent TCP MSS from being inherited from the parent, use the **tcp mss inheritance-disable** command in the appropriate configuration mode.

tcp mss inheritance-disable

Syntax Description

This command has no keywords or arguments.

Command Default

None

Command Modes

Router configuration mode

Neighbor configuration mode

Neighbor group configuration mode

Command History

Release	Modification
Release 5.3.2	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 TCP MSS under a specific neighbor:

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#router bgp 10
RP/0/RSP0/CPU0:router(config-bgp)#address-family ipv4 unicast
RP/0/RSP0/CPU0:router(config-bgp-af)#exit
RP/0/RSP0/CPU0:router(config-bgp)#neighbor-group n1
RP/0/RSP0/CPU0:router(config-bgp-nbrgrp)#tcp mss 500
RP/0/RSP0/CPU0:router(config-bgp-nbrgrp)#address-family ipv4 unicast
RP/0/RSP0/CPU0:router(config-bgp-nbrgrp-af)#exit
RP/0/RSP0/CPU0:router(config-bgp-nbrgrp)#exit
RP/0/RSP0/CPU0:router(config-bgp)#neighbor 10.0.0.2
RP/0/RSP0/CPU0:router(config-bgp-nbr)#remote-as 1
RP/0/RSP0/CPU0:router(config-bgp-nbr)#use neighbor-group n1
RP/0/RSP0/CPU0:router(config-bgp-nbr)#tcp mss inheritance-disable
```

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 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

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/RSP0/CPU0:router(config)# router bgp 109
RP/0/RSP0/CPU0:router(config-bgp)# neighbor 192.168.40.24
RP/0/RSP0/CPU0:router(config-bgp-nbr)# remote-as 1
RP/0/RSP0/CPU0:router(config-bgp-nbr)# timers 70 210
```

Related Commands

Command	Description
af-group, on page 28	Creates an address family group for BGP neighbors and enters address family group configuration mode.
neighbor-group, on page 231	Creates a neighbor group and enters neighbor group configuration mode.
session-group, on page 307	Creates a session group and enters session group configuration mode.
timers bgp, on page 569	Adjusts BGP network timers for all BGP neighbors.

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 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

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/RSP0/CPU0:router(config)# router bgp 1
RP/0/RSP0/CPU0:router(config-bgp)# timers bgp 30 90
```

Related Commands

Command	Description
timers (BGP), on page 567	Adjusts BGP network timers for a BGP neighbor.

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 4.2.1	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.

The transport can be set to either TCP or SSH. An SSH transport session with port number 22 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: once server can be TCP port 980, another can be SSH port 22, for example. This can be changed by configuration. 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:

```
Router#configure
Router(config)#router bgp 100
```

```
Router(config-bgp)#rpki server 172.168.35.40
Router(config-bgp-rpki-server)# transport ssh port 22
```

ttl-security

To configure a router to check the time-to-live (TTL) field in incoming IP packets for the specified external Border Gateway Protocol (eBGP) peer, use the **ttl-security** command in an appropriate configuration mode. To disable TTL verification, use the **no** form of this command.

ttl-security [**inheritance-disable**]
no ttl-security [**inheritance-disable**]

Syntax Description

inheritance-disable	(Optional) Prevents the **ttl-security** command from being inherited from a session group or neighbor group.

Command Default

TTL verification is not enabled for eBGP peers.

Command Modes

Neighbor configuration

VRF neighbor configuration

Neighbor group configuration

Session group configuration

Command History

Release	Modification
Release 3.7.2	This command was introduced.
Release 3.9.0	The **disable** keyword was replaced with the **inheritance-disable** keyword.

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.

Use the **ttl-security** command to enable a lightweight security mechanism to protect eBGP peering sessions from CPU utilization-based and other resource exhaustion-based attacks. These types of attacks are typically brute-force Denial of Service (DoS) attacks that attempt to disable the network by flooding devices in the network with IP packets that contain forged source and destination IP addresses in the packet headers.

This command leverages existing behavior in IP packets. For a given IP packet, the TTL count of the packet always is equal to or less than the TTL count when the packet originated, a behavior that is considered impossible to circumvent. Therefore, a packet received with a TTL count equal to the maximum TTL value of 255 can be sent only by a directly adjacent peer. When the **ttl-security** command is configured for an eBGP neighbor that is directly adjacent, the router accepts only IP packets with a TTL count that is equal to the maximum TTL value.

The **ttl-security** command secures the eBGP session in the incoming direction only. In the outbound direction, it causes packets to be sent only with the maximum TTL value so that the BGP neighbor can also verify the TTL value of incoming packets. When this command is enabled, BGP establishes or maintains a session only if the TTL value in the IP packet header is equal to the maximum TTL value. If the value is less than the maximum TTL value, the packet is discarded and an Internet Control Message Protocol (ICMP) message is not generated. This behavior is designed because a response to a forged packet is not necessary.

Note The **ttl-security** command must be configured on each participating router. Failure to configure this command on both ends of the BGP session results in the session progressing as far as the OpenSent or OpenConfirm state, remaining there until the hold time expires.

The following restrictions apply to the configuration of this command:

- The **ttl-security** command should not be configured for a peer that is already configured with the **neighbor ebgp-multihop** command. The simultaneous configuration of these commands is permitted; however, the **ttl-security** command overrides the **ebgp-multihop** command.
- This command is not supported for internal BGP (iBGP) peers.
- This command is not effective against attacks from a directly adjacent peer that has been compromised.

If you configure this command 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.

Note If the **ttl-security** command is configured on a neighbor to which the router has an established connection or the router is in the process of establishing a connection, the session must be cleared using the **clear bgp** command.

Task ID

Task ID	Operations
bgp	read, write

Examples

The following example shows how to enable TTL security for eBGP neighbor 192.168.223.7:

```
RP/0/RSP0/CPU0:router(config)# router bgp 65534
RP/0/RSP0/CPU0:router(config-bgp)# neighbor 192.168.223.7
RP/0/RSP0/CPU0:router(config-bgp-nbr)# remote-as 65507
RP/0/RSP0/CPU0:router(config-bgp-nbr)# ttl-security
```

The following example shows how to enable TTL security for multiple eBGP neighbors using a session group:

```
RP/0/RSP0/CPU0:router(config)# router bgp 65534
RP/0/RSP0/CPU0:router(config-bgp)# session-group ebgp-nbrs
RP/0/RSP0/CPU0:router(config-bgp-sngrp)# ttl-security
RP/0/RSP0/CPU0:router(config-bgp-sngrp)# exit
RP/0/RSP0/CPU0:router(config-bgp)# neighbor 192.168.223.1
RP/0/RSP0/CPU0:router(config-bgp-nbr)# remote-as 65501
RP/0/RSP0/CPU0:router(config-bgp-nbr)# use session-group ebgp-nbrs
RP/0/RSP0/CPU0:router(config-bgp-nbr)# exit
RP/0/RSP0/CPU0:router(config-bgp)# neighbor 192.168.223.2
RP/0/RSP0/CPU0:router(config-bgp-nbr)# remote-as 65502
RP/0/RSP0/CPU0:router(config-bgp-nbr)# use session-group ebgp-nbrs
RP/0/RSP0/CPU0:router(config-bgp-nbr)# exit
RP/0/RSP0/CPU0:router(config-bgp)# neighbor 192.168.223.3
```

```
RP/0/RSP0/CPU0:router(config-bgp-nbr)# remote-as 65503
RP/0/RSP0/CPU0:router(config-bgp-nbr)# use session-group ebgp-nbrs
RP/0/RSP0/CPU0:router(config-bgp-nbr)# exit
```

Related Commands

Command	Description
ebgp-multihop, on page 181	Accepts and attempts BGP connections to external peers residing on networks that are not directly connected.
neighbor-group, on page 231	Creates a neighbor group and enters neighbor group configuration mode.
session-group, on page 307	Creates a session group and enters session group configuration mode.
show lpts flows	Displays information about locally terminated packet flows, including the minimum TTL value expected.

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 router configuration 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

Router configuration

Command History

Release	Modification
Release 4.2.0	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

This example shows how to disable inbound update message basic error handling for eBGP neighbors:

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#router bgp 100
RP/0/RSP0/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/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#router bgp 100
RP/0/RSP0/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 router configuration 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

Router configuration

Command History

Release	Modification
Release 4.2.0	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

This example shows how to enable inbound update message extended error handling for eBGP neighbors:

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#router bgp 100
RP/0/RSP0/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/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#router bgp 100
RP/0/RSP0/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 Global Configuration 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

Router configuration

Command History

Release	Modification
Release 7.10.1	This command was introduced.
Release 7.11.1	This command was deprecated.

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/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#router bgp 100
RP/0/RSP0/CPU0:router(config-bgp)#neighbor 10.2.3.4
RP/0/RSP0/CPU0:router(config-bgp-nbr)#update in labeled-unicast equivalentexit
```

update limit

To set upper bound on transient memory usage for update generation, use the **update limit** command in router configuration 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

Router configuration

Command History

Release	Modification
Release 4.2.0	This command was introduced and replaced the **bgp write-limit** command.

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.

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:

```
Router(config)# router bgp 65000
Router(config-bgp)#update limit 1024
```

Related Commands

Command	Description
update limit address-family	
update limit sub-group	

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 4.2.0	This command was introduced and replaced the **bgp write-limit** command.

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.

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:

```
Router#configure
Router(config)#router bgp 100
Router(config-bgp)#address-family ipv4 unicast
```

```
Router(config-bgp-af)#update limit address-family 512
```

Related Commands

Command	Description
update limit	
update limit sub-group	

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 4.2.0	This command was introduced and replaced the **bgp write-limit** command.

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.

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:

```
Router#configure
```

```
Router(config)#router bgp 100
Router(config-bgp)#address-family ipv4 unicast
Router(config-bgp-af)#update limit sub-group ebgp 256
```

Related Commands

Command	Description
update limit	
update limit address-family	

update out logging

To enable logging of update generation events, use the **update out logging** command in router configuration 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

Router configuration

Command History

Release	Modification
Release 4.2.0	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

This example shows how to enable logging of update generation events:

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#router bgp 100
RP/0/RSP0/CPU0:router(config-bgp)#update out logging
```

update out originator-loopcheck disable

To bypass originator loop check for outgoing update, use the **update out originator-loopcheck disable** command in neighbor adress-family configuration mode. This configuration is effective only for the neighbor in unique update group for which the update is being processed.

update out originator-loopcheck [disable]

To restore the system to its default condition, use the **no** form of this command or use the **update out originator-loopcheck** command.

no update out originator-loopcheck [disable]
update out originator-loopcheck

Syntax Description

disable (optional)	Prevents the **update out originator-loopcheck disable** command from being inherited from a parent group.

Command Default

Originator loop check for outbound updates is enabled if there is only one neighbor, and disabled if there are multiple neighbors in the update group.

Command Modes

Non-VRF neighbor address family configuration

Supported Address-families

- address-family vpnv4 unicast
- address-family vpnv6 unicast
- address-family l2vpn evpn

Command History

Release	Modification
Release 6.6.2	This command was introduced.

Usage Guidelines

Use the **update out originator-loopcheck disable** command to bypass originator loop check for outgoing update. Use this command either with **as-override** command or with **as-path-loopcheck out disable** command for eBGP peering. For iBGP peering, separate configuration is not required.

The word “originator” represents the either of the following:

- The neighbor whose IP address matches with the current network path's neighbor IP address and the update is sent to that same neighbor.
- The neighbor for the route reflector or iBGP, when the neighbor's router-id matches with the current network path's router-id, and the update is sent to that same neighbor.

Although, the **disable** keyword is optional, you must configure it to disable the originator loop check. For example, under neighbor address family, if the **update out originator-loopcheck disable** command is configured, only then it disables the originator loop check for that neighbor. There can be scenarios where configuring the command without the optional **disable** keyword is useful.

Task ID

Task ID	Operation
BGP	read, write

Example 1:

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#router bgp 100
RP/0/RSP0/CPU0:router(config-bgp)# neighbor 10.40.0.1
RP/0/RSP0/CPU0:router(config-bgp)# remote-as 200
RP/0/RSP0/CPU0:router (config-bgp-nbr)#address-family l2vpn evpn
RP/0/RSP0/CPU0:router (config-bgp-nbr-af)#update out originator-loopcheck disable
RP/0/RSP0/CPU0:router(config-bgp-nbr-af)#as-override
```

Example 2:

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#router bgp 100
RP/0/RSP0/CPU0:router(config-bgp)#address-family l2vpn evpn
RP/0/RSP0/CPU0:router(config-bgp-af)#as-path-loopcheck out disable
RP/0/RSP0/CPU0:router(config-bgp)#neighbor 10.40.0.1
RP/0/RSP0/CPU0:router(config-bgp-nbr)#remote-as 200
RP/0/RSP0/CPU0:router (config-bgp-nbr)#address-family l2vpn evpn
RP/0/RSP0/CPU0:router (config-bgp-nbr-af)#update out originator-loopcheck disable
```

Example 3:

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#router bgp 100
RP/0/RSP0/CPU0:router(config-bgp)#neighbor 10.60.0.1
RP/0/RSP0/CPU0:router(config-bgp-nbr)#remote-as 100
RP/0/RSP0/CPU0:router (config-bgp-nbr)#address-family vpnv4 unicast
RP/0/RSP0/CPU0:router (config-bgp-nbr-af)#update out originator-loopcheck disable
```

Related Commands

Command	Description
as-override, on page 45	
as-path-loopcheck out disable, on page 47	
allowas-in, on page 40	

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 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

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/RSP0/CPU0:router(config)# router bgp 110
RP/0/RSP0/CPU0:router(config-bgp)# neighbor 172.20.16.6
RP/0/RSP0/CPU0:router(config-bgp-nbr)# remote-as 110
RP/0/RSP0/CPU0:router(config-bgp-nbr)# update-source Loopback0
```

Related Commands

Command	Description
neighbor-group, on page 231	Creates a neighbor group and enters neighbor group configuration mode.
session-group, on page 307	Creates a session group and enters session group configuration mode.

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

Command History

Release	Modification
Release 4.3.0	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

This example shows how to enable the **update wait-install** configuration under VPNv4 unicast address family:

```
Router#configure
Router(config)#router bgp 65500
Router(config-bgp)#address-family vpnv4 unicast
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:

Neighbor group configuration

Neighbor configuration

VRF neighbor configuration

For **use session-group** version:

Neighbor group configuration

Neighbor configuration

VRF neighbor configuration

Session-group 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

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 other neighbor 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.

```
Router(config)# router bgp 1
Router(config-bgp)# session-group session1
Router(config-bgp-sngrp)# advertisement-interval 40
Router(config-bgp-sngrp)# timers 30 90
Router(config-bgp-sngrp)# exit
Router(config-bgp)# neighbor 172.168.40.24
Router(config-bgp-nbr)# remote-as 2
Router(config-bgp-nbr)# use session-group session1
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.

```
Router(config)# router bgp 1
Router(config-bgp)# session-group session1
Router(config-bgp-sngrp)# advertisement-interval 40
Router(config-bgp-sngrp)# timers 30 90
Router(config-bgp-sngrp)# exit
Router(config-bgp)# neighbor 172.168.40.24
Router(config-bgp-nbr)# remote-as 2
Router(config-bgp-nbr)# use session-group session1
Router(config-bgp-nbr)# timers 60 180
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:

```
Router(config)# router bgp 1
Router(config-bgp)# af-group family1 address-family ipv4 multicast
```

```
Router(config-bgp-afgrp)# route-policy mcast-in in
Router(config-bgp-afgrp)# exit
Router(config-bgp)# neighbor-group neighbor1
Router(config-bgp-nbrgrp)# remote-as 2
Router(config-bgp-nbrgrp)# address-family ipv4 unicast
Router(config-bgp-nbrgrp-af)# route-policy policy1 in
Router(config-bgp-nbrgrp-af)# route-policy policy1 out
Router(config-bgp-nbrgrp-af)# exit
Router(config-bgp-nbrgrp)# address-family ipv4 multicast
Router(config-bgp-nbrgrp-af)# route-policy policy1 in
Router(config-bgp-nbrgrp-af)# route-policy policy1 out
Router(config-bgp-nbrgrp-af)# exit
Router(config-bgp)# neighbor 172.168.40.24
Router(config-bgp-nbr)# use neighbor-group neighbor1
Router(config-bgp-nbr)# address-family ipv4 multicast
Router(config-bgp-nbr-af)# use af-group family1
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:

```
Router(config)# router bgp 1
Router(config-bgp)# session-group session1
Router(config-bgp-sngrp)# advertisement-interval 40
Router(config-bgp-sngrp)# exit
Router(config-bgp)# session-group session2
Router(config-bgp-sngrp)# use session-group session1
Router(config-bgp-sngrp)# update-source Loopback0
Router(config-bgp-sngrp)# exit
Router(config-bgp)# neighbor 172.168.40.24
Router(config-bgp-nbr)# remote-as 1
Router(config-bgp-nbr)# use session-group session2
Router(config-bgp-nbr)# exit
```

Related Commands

Command	Description
af-group, on page 28	Creates an address family group for BGP neighbors and enters address family group configuration mode.
session-group, on page 307	Creates a session group and enters session group configuration mode.
neighbor-group, on page 231	Creates a neighbor group and enters neighbor group configuration mode.
remote-as (BGP), on page 277	Creates a BGP neighbor and begins the exchange of routing information.
show bgp af-group, on page 335	Displays information about BGP configuration for address family groups.

Command	Description
show bgp neighbor-group, on page 413	Displays information about the BGP configuration for neighbor groups.
show bgp neighbors, on page 417	Displays information about BGP neighbors.
show bgp session-group, on page 490	Displays information about the BGP configuration for session groups.

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 4.2.1	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.

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/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#router bgp 100
RP/0/RSP0/CPU0:router(config-bgp)#rpki server 172.168.35.40
RP/0/RSP0/CPU0:router(config-bgp-rpki-server)# transport ssh port 22
RP/0/RSP0/CPU0:router(config-bgp-rpki-server)#username rpki-user
```

vrf (BGP)

To configure a VPN routing and forwarding (VRF) instance and enter VRF configuration mode, use the **vrf** command in router configuration 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.

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

Router 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

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/RSP0/CPU0:router(config)# router bgp 1
RP/0/RSP0/CPU0:router(config-bgp)# vrf vrf-1
RP/0/RSP0/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

VRF IPv6 neighbor address family 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

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 Cisco ASR 9000 Series Router* module of the *Routing Configuration Guide for Cisco ASR 9000 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 160 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/RSP0/CPU0:router(config)# router bgp 1
RP/0/RSP0/CPU0:router(config-bgp)# neighbor 172.20.16.6
RP/0/RSP0/CPU0:router(config-bgp-nbr)# remote-as 1
RP/0/RSP0/CPU0:router(config-bgp-nbr)# address-family ipv4 unicast
RP/0/RSP0/CPU0:router(config-bgp-nbr-af)# weight 50
RP/0/RSP0/CPU0:router(config-bgp-nbr-af)# exit
```

Related Commands

Command	Description
af-group, on page 28	Creates an address family group for BGP neighbors and enters address family group configuration mode.
clear bgp, on page 137	Resets a group of BGP neighbors.
neighbor-group, on page 231	Creates a neighbor group and enters neighbor group configuration mode.
session-group, on page 307	Creates a session group and enters session group configuration mode.
set weight	Sets the weight for BGP routes.



BGP Flowspec Commands

This module provides command line interface (CLI) commands for configuring BGP Flowspec on the Cisco ASR 9000 Series Router.

- class-map type traffic (BGP-flowspec), on page 602
- class type traffic, on page 603
- destination-prefix in, on page 604
- drop (BGP-flowspec), on page 606
- flowspec, on page 607
- flowspec disable, on page 608
- local-install, on page 609
- match destination-address, on page 610
- match destination-port, on page 611
- match fragment-type, on page 612
- match icmp code, on page 613
- match icmp type, on page 614
- match packet length, on page 615
- match source-address, on page 616
- match source-port, on page 618
- match tcp flag, on page 619
- redirect (BGP Flowspec), on page 620
- service-policy, on page 621
- show flowspec, on page 622
- source prefix, on page 623

class-map type traffic (BGP-flowspec)

To define a traffic class and the associated rules that match packets to the class, use the **class-map type traffic** command inGlobal configuration mode. To remove an existing class map from the router, use the **no** form of this command.

class-map type traffic match-all *class-map-name*

Syntax Description

match-all	Specifies a match on all of the match criteria.
class-map-name	Name of the class for the class map.

Command Default

None

Command Modes

Global configuration

Command History

Release	Modification
Release 5.2.0	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.

This example shows how to specify class305 as the name of a class and defines a class map for this class.

```
RP/0/RSP0/CPU0:router# config
RP/0/RSP0/CPU0:router(config)# class-map type traffic match-all class305
RP/0/RSP0/CPU0:router(config-cmap)# match destination-address ipv4 59.2.1.2 255.255.255.0
```

class type traffic

To associate a previously configured traffic class with the policy map, and to enter the configuration mode for the specified system class, use the **class type traffic** command in the policy map configuration mode.

class type traffic *class-name*

Syntax Description

class-name	Name of the class for the class map. The class name is used for the class map and to configure policy for the class in the policy map.

Command Default

None

Command Modes

Policy map configuration mode

Command History

Release	Modification
Release 5.2.0	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.

This example shows how to associate a class map with the policy map:

```
RP/0/RSP0/CPU0:router# config
RP/0/RSP0/CPU0:router(config)# policy-map type pbr p1
RP/0/RSP0/CPU0:router(config-pmap)# class type traffic c1
RP/0/RSP0/CPU0:router(config-pmap-c)# set dscp 34
```

destination-prefix in

To filter flowspec based on destination in flowspec network-layer reachability information (NLRI) using RPL, and apply on neighbor attach point, use the **destination-prefix in** command in route-policy configuration mode.

destination-prefix in { *prefix-set-name inline-prefix-set parameter* }

Syntax Description

prefix-set-name	Name of a prefix set.
inline-prefix-set	Inline prefix set. The inline prefix set must be enclosed in parentheses.
parameter 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 5.3.2	This command was introduced.

Usage Guidelines

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

Note

- Use the **destination-prefix in** command as a conditional expression within an **if** statement.
- For a list of all conditional expressions available within an **if** statement, see the **if** command.
- This command takes either a named prefix set or an inline prefix set value as an argument. The condition returns true if the destination entry matches any entry in the prefix set or inline prefix set. An attempt to match a destination using a prefix set that is defined but contains no elements returns false.
- The routing policy language (RPL) provides the ability to test destinations for a match to a list of prefix match specifications using the **in** operator. The **destination-prefix** command is protocol-independent.
- Use the **destination-prefix in** command to filter a flowspec route in BGP, and the **match destination-address** command to create a flowspec rule that matches a specific destination.
- Use the **destination in** command to filter BGP routes for other address families, and use the **destination-prefix in** command to filter flowspec routes in BGP that include a **match destination-address** clause.
- RPL supports both 32-bit IPv4 prefixes, specified in dotted-decimal format, and 128-bit IPv6 prefixes, specified in colon-separated hexadecimal format.

Task ID

Task ID	Operations
route-policy	read, write

Examples

In this example, prefix filtering is done based on flowspec destination address:

```
RP/0/RSP0/CPU0:router(config)# route-policy policy-A
RP/0/RSP0/CPU0:router(config-rpl)# if destination-prefix in pfx then

RP/0/RSP0/CPU0:router(config-rpl-if)# drop
RP/0/RSP0/CPU0:router(config-rpl-if)# endif
RP/0/RSP0/CPU0:router(config-rpl)# end-policy
```

This example demonstrates a route policy and the point where it is applied.

```
prefix-set ipv4_flow2
150.1.1.0/24,
150.2.1.0/24
end-set
!

route-policy ipv4_dest_pass
if destination-prefix in ipv4_flow2 then
pass
else
drop
endif
end-policy
!

router bgp 100
bgp router-id 1.1.1.1
address-family ipv4 unicast
!
address-family ipv6 unicast
!
address-family ipv4 flowspec
!
address-family ipv6 flowspec
!
neighbor 33.1.1.2
remote-as 200
address-family ipv4 unicast
route-policy pass in
route-policy pass out
!
address-family ipv4 flowspec
route-policy ipv4_dest_pass in
!
!
```

drop (BGP-flowspec)

To configure a traffic class to discard packets belonging to a specific class, use the **drop** command in policy-map class configuration mode. To disable the packet discarding action in a traffic class, use the **no** form of this command.

drop
no drop

Syntax Description

This command has no keywords or arguments.

Command Default

Disabled

Command Modes

Policy-map class configuration (config-pmap-c)

Command History

Release	Modification
Release 5.2.0	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.

Examples

This example shows how to discard packets:

```
RP/0/RSP0/CPU0:router#config
RP/0/RSP0/CPU0:router(config)# policy -map type pbr match_dest_110.1.1.x_drop
RP/0/RSP0/CPU0:router(config-pmap)# class type traffic match_dest_110.1.1.x
RP/0/RSP0/CPU0:router(config-pmap-c)# drop
```

flowspec

To enter BGP flowspec configuration mode, use the **flowspec** command in Global configuration mode.

flowspec

Syntax Description

This command has no keywords or arguments.

Command Default

No default behavior or values

Command Modes

Global configuration

Command History

Release	Modification
Release 5.2.0	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.

Examples

This example show how to enter flowspec configuration mode.

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# flowspec
RP/0/RSP0/CPU0:router(config-flowspec)#
```

flowspec disable

To disable flowspec configuration on all interfaces, use the **flowspec disable** command in interface configuration mode.

ipv4 | **ipv6**
flowspec disable

Syntax Description

ipv4	Specifies IPv4 interfaces.
ipv6	Specifies IPv6 interfaces.

Command Default

No default behavior or values

Command Modes

Interface configuration

Command History

Release	Modification
Release 5.2.0	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.

Examples

This example shows how to disable flowspec configuration on all interfaces.

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# interface GigabitEthernet 0/2/0/2
RP/0/RSP0/CPU0:router(config-if)# ipv4 flowspec disable
```

local-install

To apply local installation of flowspec policy on all interfaces, use the **local-install** command in appropriate command mode.

local-install interface-all

Syntax Description

interface-all	Installs flowspec policy on all interfaces.

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 5.2.0	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.

Examples

This example show how to install flowspec policy on all interfaces under flowspec subaddress family configuration mode.

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# flowspec
RP/0/RSP0/CPU0:router(config-flowspec)# address-family ipv4
RP/0/RSP0/CPU0:router(config-flowspec-af)# local-install interface-all
```

match destination-address

To identify a specific destination IP address explicitly as a match criterion in a class map, use the **match destination-address** command in the class map configuration mode. To remove a specific destination IP address from the matching criteria for a class map, use the **no** form of this command.

match destination-address {**ipv4** | **ipv6**} *address*
no match destination-address {**ipv4** | **ipv6**} *address*

Syntax Description

ipv4	Indicates an IPv4 address.
ipv6	Indicates an IPv6 address.
address	Specifies a destination address.

Command Default

No default behavior or values

Command Modes

Class map configuration

Command History

Release	Modification
Release 5.2.0	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.

Examples

This example shows how to match a destination ipv4 address:

```
RP/0/RSP0/CPU0:router(config)#class-map type traffic match-all
RP/0/RSP0/CPU0:router(config-cmap)# match destination-address ipv4 59.2.1.2 255.255.255.0
```

match destination-port

To identify a specific destination port as the match criterion for a class map, use the **match destination-port** command in class map configuration mode. To remove destination port-based match criteria from a class map, use the **no** form of this command.

match destination-port {*destination-port-value* }
no match destination-port {*destination-port-value* }

Syntax Description

destination-port-value	A port Number. Range is from 0 to 65535.

Command Default

No default behavior or values

Command Modes

Class map configuration

Command History

Release	Modification
Release 5.2.0	The *min-value* and *max-value* variables were added.

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.

Examples

This example shows how to match a destination port:

```
RP/0/RSP0/CPU0:router(config)# class-map type traffic match-all
RP/0/RSP0/CPU0:router(config-cmap)# match destination-port 1
```

match fragment-type

To identify a fragment-type as the match criterion for a class map, use the **match fragment-type** command in class map configuration mode. To remove fragment-type match criteria from a class map, use the **no** form of this command.

match fragment type [**dont-fragment**] [**first-fragment**] [**is-fragment**] [**last-fragment**]
no match fragment type [**dont-fragment**] [**first-fragment**] [**is-fragment**] [**last-fragment**]

Syntax Description

dont-fragment	Matches dont-fragment bit.
first-fragment	Matches first-fragment bit.
is-fragment	Matches is-fragment bit.
last-fragment	Matches last-fragment bit.

Command Default

No default behavior or values

Command Modes

Class map configuration

Command History

Release	Modification
Release 5.2.0	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.

Examples

This example shows how to match a fragment-type:

```
RP/0/RSP0/CPU0:router(config)# class-map type traffic match-all
RP/0/RSP0/CPU0:router(config-cmap)# match fragment-type is-fragment
```

match icmp code

To identify an ICMP (Internet Control Message Protocol) code as the match criterion for a class map, use the **match icmp type** command in the class map configuration mode. To remove the icmp code-based match criteria from a class map, use the **no** form of this command.

match {**ipv4** | **ipv6**} **icmp-code** {*value*}
no match {**ipv4** | **ipv6**} **icmp-code** {*value*}

Syntax Description

ipv4	Indicates an IPv4 ICMP code.
ipv6	Indicates an IPv6 ICMP code.

Command Default

No default behavior or values

Command Modes

Class map configuration

Command History

Release	Modification
Release 5.2.0	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.

Examples

This example shows how to match an IPv4 ICMP code:

```
RP/0/RSP0/CPU0:router(config)# class-map type traffic match-all
RP/0/RSP0/CPU0:router(config-cmap)# match ipv4 icmp-code 1
```

match icmp type

To identify an ICMP (Internet Control Message Protocol) type as the match criterion for a class map, use the **match icmp type** command in class map configuration mode. To remove the icmp type-based match criteria from a class map, use the **no** form of this command.

match {**ipv4** | **ipv6**} **icmp-type** {*value*}
no match {**ipv4** | **ipv6**} **icmp-type** {*value*}

Syntax Description

ipv4	Indicates an IPv4 ICMP type.
ipv6	Indicates an IPv6 ICMP type.

Command Default

No default behavior or values

Command Modes

Class map configuration

Command History

Release	Modification
Release 5.2.0	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.

Examples

This example shows how to match an IPv4 ICMP type:

```
RP/0/RSP0/CPU0:router(config)# class-map type traffic match-all
RP/0/RSP0/CPU0:router(config-cmap)# match ipv4 icmp-type 1
```

match packet length

To specify the packet length in the IP header as a match criterion in a class map, use the **match packet length** command in class-map configuration mode. To remove a previously specified packet length as a match criterion, use the **no** form of this command.

match packet length {*value* }
no match packet length {*value* }

Syntax Description

value	IP packet length. Range is from 0 to 65535.

Command Default

No default behavior or values.

Command Modes

Class map configuration

Command History

Release	Modification
Release 5.2.0	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.

Examples

This example shows how to match a packet length value:

```
RP/0/RSP0/CPU0:router(config)# class-map type traffic match-all
RP/0/RSP0/CPU0:router(config-cmap)# match packet length 3
```

match source-address

To identify a specific source MAC address or an IP address explicitly as a match criterion in a class map, use the **match source-address** command in the class map configuration mode. To remove a specific source MAC address or an IP address from the matching criteria for a class map, use the **no** form of this command.

match source-address { **mac** | **ipv4** | **ipv6** } *address*
no match source-address { **mac** | **ipv4** | **ipv6** } *address*

Syntax Description

mac	Indicates a MAC address.
ipv4	Indicates an IPv4 address.
ipv6	Indicates an IPv6 address.
address	Specifies a source address.

Command Default

No default behavior or values

Command Modes

Class map configuration

Command History

Release	Modification
Release 3.7.2	This command was introduced for MAC addresses.
Release 5.2.0	Support for IP addresses was added.

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.

The **match source-address** command specifies a source address that is used as the match criterion against which packets are checked to determine if they belong to the class specified by the class map.

To use the **match source-address** command, you must first enter the **class-map** command to specify the name of the class whose match criteria you want to establish. If you specify more than one **match source-address** command in a class map, only the last command entered applies.

This command is supported on an input service policy only.

Layer 2 match criteria on a Layer 3 target, or Layer 3 match criteria on a Layer 2 target is not allowed.

The **match source-address** command is supported on egress Layer 2 interfaces, Layer 2 subinterfaces, and Layer 3 physical interfaces. Layer 3 physical interfaces are supported, because it is possible for a Layer 3 interface to have underlying Layer 2 subinterfaces.

The **match source-address** command is allowed on a policy map that is attached to an Ethernet interface. The command is invalid on a policy that is attached to a Packet-over-SONET/SDH (POS) interface or a routed VLAN subinterface.

The match 48-bit MAC address is specified in xxxx.xxxx.xxxx format on L2VPN PE interfaces.

Examples

This example shows how to match a source MAC address:

```
RP/0/RSP0/CPU0:router(config)# class-map match-any A
RP/0/RSP0/CPU0:router(config-cmap)# match source-address mac 0003.f0d0.2356
```

This example shows how to match a source IPv4 address:

```
RP/0/RSP0/CPU0:router(config)#class-map type traffic match-all A
RP/0/RSP0/CPU0:router(config-cmap)# match source-address ipv4 59.2.1.2 255.255.255.0
```

match source-port

To identify a specific source port as the match criterion for a class map, use the **match source port** command in class map configuration mode. To remove source port-based match criteria from a class map, use the **no** form of this command.

match source-port {*source-port-value* }
no match source-port {*source-port-value* }

Syntax Description

source-port-value	A port Number. Range is from 0 to 65535.

Command Default

No default behavior or values

Command Modes

Class map configuration

Command History

Release	Modification
Release 5.2.0	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.

Examples

This example shows how to match a source port:

```
RP/0/RSP0/CPU0:router(config)# class-map type traffic match-all
RP/0/RSP0/CPU0:router(config-cmap)# match source-port 1
```

match tcp flag

To identify a TCP flag as the match criterion for a class map, use the **match tcp flag** command in class map configuration mode. To remove the tcp flag based match criteria from a class map, use the **no** form of this command.

match tcp-flag *value* **any**
no match tcp-flag *value***any**

Syntax Description

value	TCP flag value. Range is from 1 to 4095 (hexadecimal).
any	Specifies a match based on any bit in the TCP flag.

Command Default

No default behavior or values

Command Modes

Class map configuration

Command History

Release	Modification
Release 5.2.0	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.

Examples

This example shows how to match a TCP flag:

```
RP/0/RSP0/CPU0:router(config)# class-map type traffic match-all
RP/0/RSP0/CPU0:router(config-cmap)# match tcp flag 2 any
```

redirect (BGP Flowspec)

To route the policy based routing (PBR) traffic to distributed denial-of-service scrubber (DDoS), use the **redirect** command in policy-map configuration mode. To return the PBR traffic to normal route, use the **no** form of this command.

redirect {**default-route** | **nexthop** } {*IPv4-address IPv6-address* | **route-target** {*AS-number: index IPv4-address: index* } | **vrf** *vrf-name*}
no redirect [default-route | nexthop]

Syntax Description

default-route	Forwards to the default nexthop for this packet
nexthop	Forwards to specified nexthop
IPv4 address	Input IPv4 Nexthop address
IPv6 address	Input IPv6 Nexthop address
route-target	Enter specific route-target string
AS-number: index	Enter 2-byte or 4-byte autonomous system number (AS) and *index* in hexa decimal or decimal format.
IPv4-address: index	Enter IPv4 address and *index* in hexa decimal or decimal format.
vrf*vrf-name*	Enter specific VRF name for the nexthop.

Command Default

None

Command Modes

Policy-map configuration

Command History

Release	Modification
Release 5.2.0	This command was introduced.

Usage Guidelines

You must be in a user group associated with a task group that includes the proper task IDs. The command reference guides include the task IDs required for each command. If you suspect user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

The example shows how to redirect PBR traffic to virtual routing and forwarding (VRF) instance:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# policy-map type pbr test1
RP/0/RSP0/CPU0:router(config-pmap)# class type traffic test1
RP/0/RSP0/CPU0:router(config-pmap-c)# redirect nexthot vrf vrf1
```

service-policy

To configure service policy on a flowspec subaddress family interface, use the **service-policy** command in appropriate command mode.

service-policy type pbr *policy-name*

Syntax Description

type	Specifies type of the service policy.
pbr	Specifies a policy-based routing (PBR) policy map.
policy-name	Name of the policy map.

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 5.2.0	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.

Examples

This example shows how to setup service policy.

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# flowspec
RP/0/RSP0/CPU0:router(config-flowspec)# address-family ipv4
RP/0/RSP0/CPU0:router(config-flowspec-af)# service-policy type pbr policy100
```

show flowspec

To display flowspec policy information for an interface, use the **show flowspec** command in EXEC mode.

show flowspec {**afi-all** | **client** | **ipv4** | **ipv6** | **summary** | **vrf**}

Syntax Description

afi-all	Displays flowspec policy applied on IPv4 and IPv6 interfaces.
client	Displays flowspec client interfaces.
ipv4	Displays flowspec policy applied on IPv4 interfaces.
ipv6	Displays flowspec policy applied on IPv6 interfaces.
summary	Displays flowspec policy summary on all interfaces.
vrf	Displays flowspec policy applied on VRF interfaces.

Command Default

No default behavior or values

Command Modes

EXEC

Command History

Release	Modification
Release 5.2.0	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.

Examples

This example shows sample output from **show flowspec** command when **vrf**, **ipv4** and **summary** keywords are used.

```
RP/0/RSP0/CPU0:router# show flowspec vrf vrf1 ipv4 summary
Mon May 19 12:59:41.226 PDT
Flowspec VRF+AFI table summary:
VRF: vrf1
  AFI: IPv4
    Total Flows:              3
    Total Service Policies:   1
```

source prefix

To filter flowspec based on source in flowspec network-layer reachability information (NLRI) using RPL, and apply on neighbor attach point, use the **source prefix** command in route-policy configuration mode.

source prefix {*prefix-set-nameinline-prefix-setparameter*}

Syntax Description

prefix-set-name	Name of a prefix set.
inline-prefix-set	Inline prefix set. The inline prefix set must be enclosed in parentheses.
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 5.3.2	This command was introduced.

Usage Guidelines

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

Use the **source prefix** command as a conditional expression within an **if** statement. A comparison that references a prefix set with zero elements in it returns false.

Note

- For a list of all conditional expressions available within an **if** statement, see the **if** command.
- The source of a BGP route is the IP peering address of the neighboring router from which the route was received.
- The prefix set can contain both IPv4 and IPv6 prefix specifications.

Task ID

Task ID	Operations
route-policy	read, write

Examples

In this example, prefix filtering is done based on flowspec source address:

```
RP/0/RSP0/CPU0:router(config)# route-policy policy-A
RP/0/RSP0/CPU0:router(config-rpl)# If source-prefix in my-prefix-set then
pass
```

Related Commands

Command	Description
prefix-set, on page 1714	Enters a prefix set configuration mode and defines a prefix set.



BFD Commands

This module provides command line interface (CLI) commands for configuring Bidirectional Forwarding Detection (BFD) on the Cisco ASR 9000 Series Router.

- address-family ipv4 unicast (BFD), on page 627
- bfd, on page 629
- bfd address-family ipv4 destination, on page 631
- bfd address-family ipv4 echo minimum-interval, on page 632
- bfd address-family ipv4 fast-detect, on page 634
- bfd address-family ipv4 minimum-interval, on page 635
- bfd address-family ipv4 multiplier, on page 638
- bfd address-family ipv4 timers, on page 641
- bundle coexistence bob-blb, on page 643
- bfd dampening, on page 645
- bfd dampening disable, on page 648
- bfd echo ipv4 bundle-per-member minimum-interval, on page 649
- bfd encap-mode, on page 651
- bfd fast-detect, on page 652
- bfd minimum-interval, on page 655
- bfd mode, on page 659
- bfd multipath include location, on page 660
- bfd multiplier, on page 662
- bfd bundle per-member mode, on page 665
- clear bfd counters, on page 666
- clear bfd dampening, on page 668
- echo disable, on page 670
- echo ipv4 source, on page 672
- echo latency detect, on page 674
- echo startup validate, on page 676
- hw-module bfd-hw-offload, on page 678
- interface (BFD), on page 680
- ipv6 checksum, on page 683
- multihop ttl-drop-threshold, on page 685
- show bfd, on page 687
- show bfd client, on page 690

- show bfd counters, on page 692
- show bfd hw-offload, on page 695
- show bfd mib session, on page 697
- show bfd multipath, on page 700
- show bfd session, on page 702
- show bfd summary, on page 708

address-family ipv4 unicast (BFD)

To enable Bidirectional Forwarding Detection (BFD) fast-detection on a specific IPV4 unicast destination address prefix and on the forwarding next-hop address, use the **address-family ipv4 unicast** command in static route configuration mode. To return the router to the default setting, use the **no** form of this command.

address-family ipv4 unicast *address nexthop* **bfd fast-detect** [**minimum interval** *interval*] [**multiplier** *multiplier*]
no address-family ipv4 unicast *address nexthop* **bfd fast-detect** [**minimum interval** *interval*] [**multiplier** *multiplier*]

Syntax Description

address	Specifies the IPv4 unicast destination address and prefix on which to enable BFD fast-detection.
nexthop	Specifies the next-hop address on which to enable BFD fast-detection.
bfd fast-detect	Enables BFD fast-detection on the specified IPV4 unicast destination address prefix and on the forwarding next-hop address.
minimum interval *interval*	(Optional) Ensures that the next hop is assigned with the same hello interval. Replace *interval* with a number that specifies the interval in milliseconds. Range is from 10 through 10000.
multiplier *multiplier*	(Optional) Ensures that the next hop is assigned with the same detect multiplier. Replace *multiplier* with a number that specifies the detect multiplier. Range is from 1 through 10.

Command Default

interval: 100

multiplier: 3

Command Modes

Static route configuration mode

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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

If the multiplier is changed using the **bfd multiplier** command, the new parameter is used to update all existing BFD sessions for the protocol (BGP, IS-IS, MPLS-TE, or OSPF).

Task ID

Task ID	Operations
static	read, write

Examples

The following example shows how to enable BFD on a static route. In this example, BFD sessions are established with the next-hop 3.3.3.3 when it becomes reachable.

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router static
RP/0/RSP0/CPU0:router (config-static)# address-family ipv4 unicast 2.2.2.0/24 3.3.3.3 bfd
fast-detection
```

Related Commands

Command	Description
bfd fast-detect, on page 652	Enables BFD to detect failures in the path between adjacent forwarding engines.
show bfd, on page 687	Displays BFD information for a specific location.

bfd

To enter Bidirectional Forwarding Detection (BFD) configuration mode, use the **bfd** command in global configuration mode. To exit BFD configuration mode and return to global configuration mode, use the **no** form of this command.

bfd
no bfd

Syntax Description

This command has no keywords or arguments.

Command Default

No default behavior or values

Command Modes

Global 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

When you issue the **bfd** command in global configuration mode, the CLI prompt changes to "config-bfd," indicating that you have entered BFD configuration mode. In the following sample output, the question mark (**?**) online help function displays all the commands available under BFD configuration mode:

```
RP/0/RSP0/CPU0:router(config)# bfd
RP/0/RSP0/CPU0:router(config-bfd)# ?

  commit     Commit the configuration changes to running
  describe   Describe a command without taking real actions
  do         Run an exec command
  echo       Configure BFD echo parameters
  exit       Exit from this submode
  interface  Configure BFD on an interface
  no         Negate a command or set its defaults
  root       Exit to the global
      configuration  mode
  show       Show contents of configuration
```

Task ID

Task ID	Operations
bgp	read, write
ospf	read, write

Task ID	Operations
isis	read, write
mpls-te	read, write

Examples

The following example shows how to enter BFD configuration mode:

```
RP/0/RSP0/CPU0:router # configure
RP/0/RSP0/CPU0:router(config)# bfd
RP/0/RSP0/CPU0:router(config-bfd)#
```

Related Commands

Command	Description
echo disable, on page 670	Disables echo mode on a router or on an individual interface or bundle.
interface (BFD), on page 680	Enters BFD interface configuration mode, where you can disable echo mode on an interface.
show bfd, on page 687	Displays BFD information for a specific location.

bfd address-family ipv4 destination

To specify the destination address for BFD sessions on bundle member links, use the **bfd address-family ipv4 destination** command in interface configuration mode. To return to the default, use the **no** form of this command.

bfd address-family ipv4 destination *ip-address*
no bfd address-family ipv4 destination *ip-address*

Syntax Description

ip-address	32-bit IPv4 address in dotted-decimal format (A.B.C.D).

Command Default

No destination IPv4 address is configured.

Command Modes

Interface configuration (config-if)

Command History

Release	Modification
Release 4.0.0	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.

This command is supported on bundle interfaces only.

Task ID

Task ID	Operations
bundle	read, write

Examples

The following example specifies the IPv4 address of 10.20.20.1 as the destination address for the BFD session on an Ethernet bundle interface:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# interface Bundle-Ether 1
RP/0/RSP0/CPU0:router(config-if)# bfd address-family ipv4 destination 10.20.20.1
```

Related Commands

Command	Description
bfd address-family ipv4 fast-detect, on page 634	Enables IPv4 BFD sessions on bundle member links.

bfd address-family ipv4 echo minimum-interval

To specify the minimum interval for echo packets on IPv4 BFD sessions on bundle member links, use the **bfd address-family ipv4 echo minimum-interval** command in interface configuration mode. To return to the default, use the **no** form of this command.

Note The interface configuration of the echo timer value takes precedence over the global echo timer value.

bfd address-family ipv4 echo minimum-interval *milliseconds*
no bfd address-family ipv4 echo minimum-interval [*milliseconds*]

Syntax Description

milliseconds	Shortest interval between sending BFD echo packets to a neighbor. The range is 15 to 2000 milliseconds.

Command Default

The default value is the product of the async and multiplier values.

Command Modes

Interface configuration

Command History

Release	Modification
Release 5.3.0	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.

This command is supported on only on BFD over Bundle Per-Member Link Mode using Cisco standard.

This command allows to configure an echo interval to a value greater than async x bfd multiplier value.

Task ID

Task ID	Operations
bundle	read, write

Examples

The following example specifies that echo packets will be sent at a minimum interval of 900 ms for IPv4 BFD sessions on member links of an Ethernet bundle:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# interface Bundle-Ether 1
RP/0/RSP0/CPU0:router(config-if)# bfd address-family ipv4 minimum-interval 200
RP/0/RSP0/CPU0:router(config-if)# bfd address-family ipv4 echo minimum-interval
 900
```

Related Commands

Command	Description
bfd address-family ipv4 multiplier, on page 638	Specifies a number that is used as a multiplier with the minimum interval to determine BFD control and echo packet failure detection times and echo packet transmission intervals for IPv4 BFD sessions on bundle member links.
bfd address-family ipv4 minimum-interval, on page 635	Specifies the minimum interval for asynchronous mode control packets on IPv4 BFD sessions on bundle member links.

bfd address-family ipv4 fast-detect

To enable IPv4 BFD sessions on bundle member links, use the **bfd address-family ipv4 fast-detect** command in interface configuration mode. To return to the default, use the **no** form of this command.

bfd address-family ipv4 fast-detect
no bfd address-family ipv4 fast-detect

Syntax Description

This command has no keywords or arguments.

Command Default

BFD sessions are disabled.

Command Modes

Interface configuration (config-if)

Command History

Release	Modification
Release 4.0.0	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.

This command is supported on bundle interfaces only.

Task ID

Task ID	Operations
bundle	read, write

Examples

The following example enables IPv4 BFD sessions on member links of an Ethernet bundle:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# interface Bundle-Ether 1
RP/0/RSP0/CPU0:router(config-if)# bfd address-family ipv4 fast-detect
```

Related Commands

Command	Description
bfd address-family ipv4 destination, on page 631	Specifies the destination address for BFD sessions on bundle member links.

bfd address-family ipv4 minimum-interval

To specify the minimum interval for asynchronous mode control packets on IPv4 BFD sessions on bundle member links, use the **bfd address-family ipv4 minimum-interval** command in interface configuration mode. To return to the default, use the **no** form of this command.

bfd address-family ipv4 minimum-interval *milliseconds*
no bfd address-family ipv4 minimum-interval [*milliseconds*]

Syntax Description

milliseconds	Shortest interval between sending BFD control packets to a neighbor. The range is 15 to 30000 milliseconds. **Note** Although the command allows you to configure a minimum of 15 ms, the supported minimum is 50 ms.

Command Default

The default is 150 ms.

Command Modes

Interface configuration (config-if)

Command History

Release	Modification
Release 4.0.0	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.

This command is supported on bundle interfaces only.

The BFD minimum interval is used with a configurable multiplier (**bfd address-family ipv4 multiplier** command) to determine the intervals and failure detection times for both control and echo packets in asynchronous mode on bundle member links.

For example, with a session interval of *I* and a multiplier of *M*, the following packet intervals and failure detection times apply for BFD asynchronous mode:

- Value of *I*—Minimum period between sending of BFD control packets.
- Value of *I* x *M*
  - BFD control packet failure detection time. This is the maximum amount of time that can elapse without receipt of a BFD control packet before the session is declared down.
  - Minimum period between sending of BFD echo packets.
- Value of (*I* x *M*) x *M*—BFD echo packet failure detection time. This is the maximum amount of time that can elapse without receipt of a BFD echo packet before the session is declared down.

When used with bundled VLANs, the following restrictions apply:

- The command specifies control packet intervals only because echo packets are not supported.
- The minimum interval is 250 ms.

The **bfd address-family ipv4 minimum-interval** command in bundle interface configuration overrides the minimum intervals specified by the **bfd minimum-interval** command in other areas of BFD configuration.

Note When multiple applications share the same BFD session, the application with the most aggressive timer is used locally. Then, the result is negotiated with the peer router.

Keep the following router-specific rules in mind when configuring the minimum BFD interval:

- The maximum rate in packets-per-second (pps) for BFD sessions is linecard-dependent. If you have multiple linecards supporting BFD, then the maximum rate for BFD sessions per system is the supported linecard rate multiplied by the number of linecards.
  - The maximum rate for BFD sessions per linecard is 9600 pps.
- The maximum number of all BFD sessions on the router is 1024.
- The maximum number of all BFD sessions on the router is 1440.

To calculate the rate for BFD sessions on bundle members running in asynchronous mode without echo:

- Divide 1000 by the value of the minimum interval (as specified by the **bfd address-family ipv4 minimum-interval** command). This is also the base rate used per member session with echo:

  Asynchronous rate per bundle member = (1000 / Min-interval)

To calculate the rate for BFD sessions on bundle members running in asynchronous mode with echo:

- Determine the echo interval, which is the value of the minimum interval (specified by the **bfd address-family ipv4 minimum-interval** command) multiplied by the multiplier value (specified by the **bfd address-family ipv4 multiplier** command).

  Echo interval = (Min-interval x Multiplier)

- Calculate the overall rate supported for all members on the bundle:

  Ethernet bundle rate = (1000 / Echo interval) x 64

- Add the asynchronous base rate per bundle member to find the total rate for all bundle links:

  Total bundle rate = Ethernet bundle rate + (Base asynchronous rate x Number of links)

Task ID

Task ID	Operations
bundle	read, write

Examples

The following example specifies that control packets will be sent at a minimum interval of 200 ms for IPv4 BFD sessions on member links of an Ethernet bundle:

```
RP/0/RSP0/CPU0:router# configure
```

```
RP/0/RSP0/CPU0:router(config)# interface Bundle-Ether 1
RP/0/RSP0/CPU0:router(config-if)# bfd address-family ipv4 minimum-interval 200
```

Related Commands

Command	Description
bfd minimum-interval, on page 655	Specifies the minimum control packet interval for BFD sessions for the corresponding BFD configuration scope.
bfd address-family ipv4 multiplier, on page 638	Specifies a number that is used as a multiplier with the minimum interval to determine BFD control and echo packet failure detection times and echo packet transmission intervals for IPv4 BFD sessions on bundle member links.

bfd address-family ipv4 multiplier

To specify a number that is used as a multiplier with the minimum interval to determine BFD control and echo packet failure detection times and echo packet transmission intervals for IPv4 BFD sessions on bundle member links, use the **bfd address-family ipv4 multiplier** command in interface configuration mode. To return to the default, use the **no** form of this command.

bfd address-family ipv4 multiplier *multiplier*
no bfd address-family ipv4 multiplier [*multiplier*]

Syntax Description

multiplier Number from 2 to 50.

Note
Although the command allows you to configure a minimum of 2, the supported minimum is 3.

Command Default

The default multiplier is 3.

Command Modes

Interface configuration (config-if)

Command History

Release	Modification
Release 4.0.0	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.

This command is supported on bundle interfaces only.

The BFD multiplier is used with a configurable minimum interval (**bfd address-family ipv4 minimum-interval** command) to determine the intervals and failure detection times for both control and echo packets in asynchronous mode on bundle member links.

For example, with a session interval of *I* and a multiplier of *M*, the following packet intervals and failure detection times apply for BFD asynchronous mode:

- Value of *I*—Minimum period between sending of BFD control packets.
- Value of *I* x *M*
  - BFD control packet failure detection time. This is the maximum amount of time that can elapse without receipt of a BFD control packet before the session is declared down.
  - Minimum period between sending of BFD echo packets.



Note The maximum echo packet interval for BFD on bundle member links is the minimum of either 30 seconds or the asynchronous control packet failure detection time.

- Value of (*I* x *M*) x *M*—BFD echo packet failure detection time. This is the maximum amount of time that can elapse without receipt of a BFD echo packet before the session is declared down.

Keep the following router-specific rules in mind when configuring the minimum BFD interval:

- The maximum rate in packets-per-second (pps) for BFD sessions is linecard-dependent. If you have multiple linecards supporting BFD, then the maximum rate for BFD sessions per system is the supported linecard rate multiplied by the number of linecards.
  - The maximum rate for BFD sessions per linecard is 9600 pps.
- The maximum number of all BFD sessions per linecard is 1024.
- The maximum number of all BFD sessions per linecard is 1440.

To calculate the rate for BFD sessions on bundle members running in asynchronous mode without echo:

- Divide 1000 by the value of the minimum interval (as specified by the **bfd address-family ipv4 minimum-interval** command). This is also the base rate used per member session with echo:

  Asynchronous rate per bundle member = (1000 / Min-interval)

To calculate the rate for BFD sessions on bundle members running in asynchronous mode with echo:

- Determine the echo interval, which is the value of the minimum interval (specified by the **bfd address-family ipv4 minimum-interval** command) multiplied by the multiplier value (specified by the **bfd address-family ipv4 multiplier** command).

  Echo interval = (Min-interval x Multiplier)

- Calculate the overall rate supported for all members on the bundle:

  Ethernet bundle rate = (1000 / Echo interval) x 64

- Add the asynchronous base rate per bundle member to find the total rate for all bundle links:

  Total bundle rate = Ethernet bundle rate + (Base asynchronous rate x Number of links)

Task ID

Task ID	Operations
bundle	read, write

Task ID Examples

The following example specifies the following packet intervals and failure detection times for IPv4 BFD sessions on member links with asynchronous echo mode on an Ethernet bundle:

- 200 ms control packet interval
- 600 ms control packet failure detection interval
- 600 ms echo packet interval
- 1800 ms echo packet failure detection interval

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# interface Bundle-Ether 1
```

```
RP/0/RSP0/CPU0:router(config-if)# bfd address-family ipv4 minimum-interval 200
RP/0/RSP0/CPU0:router(config-if)# bfd address-family ipv4 multiplier 3
```

Related Commands

Command	Description
bfd address-family ipv4 minimum-interval, on page 635	Specifies the minimum interval for asynchronous mode control packets on IPv4 BFD sessions on bundle member links.
bfd minimum-interval, on page 655	Specifies the minimum control packet interval for BFD sessions for the corresponding BFD configuration scope.

bfd address-family ipv4 timers

To configure timers to allow for delays in receipt of BFD state change notifications (SCNs) from peers before declaring a link bundle BFD session down for IPv4 BFD sessions on bundle member links, use the **bfd address-family ipv4 timers** command in interface configuration mode. To return to the default, use the **no** form of this command.

Note BFD session flaps when the **show tech-support { bfd | cef | spp }** command is executed on the system with 30 millisecond timer configured. This occurs only when Cisco ASR 9000 Series Ethernet Line Card is installed in the system.

bfd address-family ipv4 timers [**start** | **nbr-unconfig**] *seconds*
no bfd address-family ipv4 timers [**start** | **nbr-unconfig**] *seconds*

Syntax Description

start *seconds*	Number of seconds after startup of a BFD member link session to wait for the expected notification from the BFD peer to be received, so that the session can be declared up. If the SCN is not received after that period of time, the BFD session is declared down. The range is 60 to 3600. **Note** In Cisco IOS XR Releases 4.0 and 4.0.1, the available minimum is 30, but is not recommended.
nbr-unconfig *seconds*	Number of seconds to wait after receipt of notification that the BFD configuration has been removed by a BFD neighbor, so that any configuration inconsistency between the BFD peers can be fixed. If the BFD configuration issue is not resolved before the specified timer is reached, the BFD session is declared down. The range is 60 to 3600. **Note** In Cisco IOS XR Releases 4.0 and 4.0.1, the available minimum is 30, but is not recommended.

Command Default

No timers are configured.

Command Modes

Interface configuration (config-if)

Command History

Release	Modification
Release 4.0.0	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.

This command is supported on bundle interfaces only.

Task ID

Task ID	Operations
bundle	read, write

Examples

The following example configures a timer for members of the specified Ethernet bundle that allows up to 1 minute (60 seconds) after startup of a BFD member link session to wait for receipt of the expected notification from the BFD peer to declare the session up:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# interface Bundle-Ether 1
RP/0/RSP0/CPU0:router(config-if)# bfd address-family ipv4 timers start 60
```

The following example configures a timer for members of the specified Ethernet bundle that allows up to 1 minute (60 seconds) to wait after receipt of notification that the BFD configuration has been removed by a BFD neighbor, before declaring a BFD session down:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# interface Bundle-Ether 1
RP/0/RSP0/CPU0:router(config-if)# bfd address-family ipv4 timers nbr-unconfig 60
```

bundle coexistence bob-blb

To configure a coexistence mechanism between BFD over Bundle (BoB) and BFD over Logical Bundle (BLB), use the **bundle coexistence bob-blb** command in BFD configuration mode. To disable this feature use the **no** form of this command.

bundle coexistence bob-blb {**inherit** | **logical**}
no bundle coexistence bob-blb {**inherit** | **logical**}

Syntax Description

inherit	Specifies inheritance as the coexistence mechanism.
logical	Specifies to use BFD logical bundle natively for coexistence.

Command Default

No coexistence configuration is enabled. The BLB configuration behaves the same as the BVLAN configuration.

Command Modes

BFD configuration

Command History

Release	Modification
Release 4.3.0	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.

The default behavior for BLB without any coexistence configuration behaves exactly as the BVLAN configuration.

When the "inherited" coexistence mode is configured, the BLB session always uses the inherited mode. It always creates a virtual session, and never a BFD session with real packets. This mode prevents the automatic conversion between BLB and the inherited mode, which could happen in default mode, depending on the sequence in which the configuration is applied. The "inherited" session is in "Down" state when BoB is not enabled.

The "logical" mode always creates BFD session with real packets.

Task ID

Task ID	Operation
bgp	read, write
ospf	read, write
isis	read, write
mpls-te	read, write

This example shows how to configure the coexistence mechanism between BFD over Bundle (BoB) and BFD over Logical Bundle (BLB) as being "inherited":

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#bfd
RP/0/RSP0/CPU0:router(config-bfd)#bundle coexistence bob-blb inherit
```

This example shows how to configure the coexitence mechanism between BFD over Bundle (BoB) and BFD over Logical Bundle (BLB) as being "logical":

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#bfd
RP/0/RSP0/CPU0:router(config-bfd)#bundle coexistence bob-blb logical
```

bfd dampening

To configure a device to dampen a flapping Bidirectional Forwarding Detection (BFD) session, use the **bfd dampening** command in global configuration mode. To restore the values of the dampening timers to their default values, use the **no** form of this command.

bfd dampening [**bundle-member**] {**initial-wait** | **l3-only-mode** | **maximum-wait** | **secondary-wait** | **extension down-monitoring** | **threshold**} *milliseconds*
no bfd dampening [**bundle-member**] {**initial-wait** | **l3-only-mode** | **maximum-wait** | **secondary-wait** | **extension down-monitoring** | **threshold**} *milliseconds*

Syntax Description

bundle-member	Specifies initial, maximum, or secondary delays in milliseconds for BFD session startup on BFD bundle members.
initial-wait	Specifies the initial delay in milliseconds before starting a BFD session. For bundle members, the default is 16000. For non-bundle interfaces, the default is 2000.
l3-only-mode	Applies dampening, if the failure is Layer 3 (L3)-specific.
maximum-wait	Specifies the maximum delay in milliseconds before starting a BFD session. **Note** The maximum delay must be greater than the initial delay.
secondary-wait	Specifies a secondary delay in milliseconds before starting a BFD session. For bundle members, the default is 20000. For non-bundle interfaces, the default is 5000.
extensions down-monitoring	Enables session monitoring extensions in down state.
threshold	Specifies a stability threshold to enable dampening in milliseconds. Range 6000 to 360000. The default is 120000.
milliseconds	For bundle members, the range is 1-518400000. For non-bundle interfaces, the range is 1-3600000.

Command Default BFD dampening is enabled by default.

Command Modes Global configuration

Command History

Release	Modification
Release 3.9.0	This command was introduced.
Release 4.0.0	The **bundle-member** keyword was added.
Release 4.2.0	The **l3-only-mode** keyword was added to the other dampening options.
Release 5.1	The **extensions** and **down-monitoring** keywords were added.

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.

Note The **initial-wait**, **maximum-wait**, and **secondary-wait** keywords are dampening timers.

You do not have to configure the BFD startup timers. If you do configure the initial wait startup timer (using the **initial-wait** keyword), then it must be less than the value of the maximum-wait timer.

By default, BFD dampening is applied to all sessions in the following manner:

- If a session is brought down, then dampening is applied before a session is allowed to transition to initial/up states.
- Length of time a session is dampened grows exponentially with continuous session flap.
- If a session remains up for minimum two minutes, then the length of time a session dampens with the next session flap is reset to the initial dampening value.

BFD on bundle member applies dampening, only if the detected failure is specific to layer 3. BFD dampening is not invoked for L1 or L2 failures. BFD is started after Layer 1 and Layer 2 (LACP) is up to prevent a race condition and false triggers. BFD is notified to stop/ignore when L1 or L2 goes down and must be notified to start/resume when L1 or L2 recovers for a given/affected link/member.

BFD applies dampening till the session transitions from up to down state and the session is not removed. Whenever there is a failure detected at L1 or L2 , the bundle manager removes BFD session on a member.

When dampening is removed a syslog message ‘Exponential backoff dampening for BFD session has been cleared for specified BFD session. When/if same session gets created by application(s), only calculated initial wait time will be applied’ is generated. If this is the desired behavior, then dampening can be enabled by configuring the BFD configuration, by using the command **bfd dampening bundle-member l3-failure-only**.

Task ID

Task ID	Operation
bgp	read, write
ospf	read, write

Task ID	Operation
isis	read, write
mpls-te	read, write

This example shows how to configure BFD dampening by specifying an initial and maximum delay for BFD session startup on BFD bundle members:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# bfd dampening bundle-member initial-wait 1000
RP/0/RSP0/CPU0:router(config)# bfd dampening bundle-member maximum-wait 3000
```

This example shows how to configure BFD dampening on a non-bundle interface:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# bfd dampening initial-wait 1000
RP/0/RSP0/CPU0:router(config)# bfd dampening maximum-wait 3000
```

bfd dampening disable

To disable a device from dampening a flapping Bidirectional Forwarding Detection (BFD) session, use the **bfd dampening disable** command in global configuration mode. To revoke the dampening of a flapping BFD session **no** form of this command.

bfd dampening disable
no bfd dampening disable

Syntax Description

This command has no arguments or keywords.

Command Default

BFD dampening is enabled by default.

Command Modes

Global configuration

Command History

Release	Modification
Release 6.1.x	This command was introduced.

Task ID

Task ID	Operation
bgp	read, write
ospf	read, write
isis	read, write
mpls-te	read, write

This example shows how to disable BFD dampening

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# bfd dampening disable
```

bfd echo ipv4 bundle-per-member minimum-interval

To globally specify the minimum global interface configuration mode. To return to the default, use the **no** form of this command.

bfd echo ipv4 bundle-per-member minimum-interval *milliseconds*
bfd echo ipv4 bundle-per-member minimum-interval [*milliseconds*]

Syntax Description

milliseconds	Shortest interval between sending BFD echo packets to a neighbor. The range is 15 to 2000 milliseconds.

Command Default

The default value is the product of the async and multiplier values.

Command Modes

Global configuration

Command History

Release	Modification
Release 5.3.0	This command was introduced.

Usage Guidelines

This command allows you to globally configure an echo interval to a value greater than async*bfd multiplier value. When the value of echo configured is lesser than the I*M1, where I is the Minimum period between sending of BFD Async packets and M1 is the remote multiplier, then the echo value is taken as I*M1 and echo detection time is taken as I*M1*M2 where M2 is the local Multiplier.

This command is supported on only on BFD over Bundle Per-Member Link mode using Cisco standard.

Task ID

Task ID	Operations
bundle	read, write

Examples

The following example specifies that echo packets will be sent globally at a minimum interval of 500 ms for IPv4 BFD sessions on bundle member links:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)#  bfd echo ipv4 bundle-per-member minimum-interval
 500
```

Related Commands

Command	Description
bfd address-family ipv4 multiplier, on page 638	Specifies a number that is used as a multiplier with the minimum interval to determine BFD control and echo packet failure detection times and echo packet transmission intervals for IPv4 BFD sessions on bundle member links.

Command	Description
bfd address-family ipv4 minimum-interval, on page 635	Specifies the minimum interval for asynchronous mode control packets on IPv4 BFD sessions on bundle member links.
bfd address-family ipv4 echo minimum-interval, on page 632	Specifies the minimum interval for asynchronous mode control packets on IPv4 BFD sessions on bundle member links.

bfd encap-mode

To enable continuity check for MPLS LSPs via BFD over Generic Associated Channel Label (GAL) (Label 13), use the **bfd encap-mode** command in MPLS Traffic Engineering Tunnel (TE) interface configuration mode. To disable continuity check for MPLS LSPs via BFD running over GAL channel, use the **no** form of this command.

bfd encap-mode gal

no bfd encap-mode gal

Syntax Description

gal	Specifies the use of BFD over Generic Associated Channel Label (GAL) (Label 13) for MPLS LSPs.

Command Default

No default behavior or values.

Command Modes

MPLS TE Tunnel interface configuration

Command History

Release	Modification
Release 5.2.0	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.

By default, BFD is enabled over an IP channel. This command enables BFD over GAL channel.

Note BFD over Generic Associated Channel Label (GAL) is not supported on virtual interfaces.

Task ID

Task ID	Operations
mpls-te	read, write

Examples

This example shows how to enable continuity check for MPLS LSPs via BFD over GAL channel:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# interface tunnel-te 1
RP/0/RSP0/CPU0:router(config-if)# bfd
RP/0/RSP0/CPU0:router(config-if-tunte-bfd)# encap-mode gal
```

bfd fast-detect

To enable Bidirectional Forwarding Detection (BFD) to detect failures in the path between adjacent forwarding engines, use the **bfd fast-detect** command in the appropriate configuration mode. To return the software to the default state in which BFD is not enabled, use the **no** form of this command.

bfd fast-detect [**disable** | **ipv4**]
no bfd fast-detect

Syntax Description

disable Disables the detection of failures in the path between adjacent forwarding engines for a specified entity, such as a BGP neighbor or OSPF interface.

Note
The **disable** keyword is available in the following modes only: BGP configuration, OSPF area configuration, OSPF area interface configuration, OSPFv3 area configuration, and OSPFv3 area interface configuration .

ipv4 Enables Intermediate System-to-Intermediate System (IS-IS) BFD detection of failures in the path between adjacent forwarding engines.

Note
The **ipv4** keyword is available in IS-IS router configuration mode only.

Command Default

BFD detection of failures in the path between adjacent forwarding engines is disabled.

Command Modes

Neighbor configuration

Session group configuration

Neighbor group configuration

Interface configuration

Interface configuration

Router configuration

Area configuration

Area interface configuration

Router configuration

Area configuration

Area interface configuration

Interface configuration

Command History

Release	Modification
Release 3.7.2	This command was introduced.

Release	Modification
Release 4.0.0	The **bfd fast-detect** command was supported in the following modes: OSPFv3 router configuration, OSPFv3 area configuration, and OSPFv3 area interface configuration.

Usage Guidelines

Note BFD can support multihop for internal and external BGP peers.

Use the **bfd fast-detect** command to provide protocol- and media-independent, short-duration failure detection of the path between adjacent forwarding engines, including the interfaces and data links.

BFD must be configured on directly connected neighbors for a BFD session to be established between the neighbors.

When MPLS-TE tunnels are protected by backup tunnels, BFD failure triggers fast reroute on affected tunnels.

In OSPF and OSPFv3 environments, the setting of the **bfd fast-detect** command is inherited from the highest-level configuration mode in which the command was configured. From the lowest to the highest configuration modes, the inheritance rules are as follows:

- If you enable BFD in area interface configuration mode, it is enabled on the specified interface only.
- If you enable BFD in area configuration mode, it is enabled on all interfaces in the specified area.
- If you enable BFD in router configuration mode, it is enabled on all areas and all associated interfaces in the specified routing process.

The **disable** keyword is available in the following modes: BGP configuration, OSPF area configuration, OSPF area interface configuration, OSPFv3 area configuration, and OSPFv3 area interface configuration . In OSPF and OSPFv3 environments, the **disable** option enables you to override the inheritance rules described previously. For example, if you enable BFD in an OSPF area, BFD is enabled on all interfaces in that area. If you do not want BFD running on one of the interfaces in that area, you must specify the **bfd fast-detect disable** command for that interface only.

To disable BFD or return the software to the default state in which BFD is not enabled in IS-IS router configuration mode and MPLS-TE configuration mode, you must enter the **no bfd fast-detect** command.

Task ID

Task ID	Operations
bgp	read, write
isis	read, write
mpls-te	read, write
ospf	read, write
multicast	read, write

Examples

The following example shows how to configure BFD on a BGP router:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router bgp 65000
RP/0/RSP0/CPU0:router(config-bgp)# neighbor 192.168.70.24
RP/0/RSP0/CPU0:router(config-bgp-nbr)# remote-as 2
RP/0/RSP0/CPU0:router(config-bgp-nbr)# bfd fast-detect
```

The following example shows the configuration of an OSPFv3 routing process named san_jose. The example shows two areas, each of which includes Gigabit Ethernet interfaces. In area 0, BFD is enabled at the area level, which means that by virtue of the inheritance rules, BFD is enabled on all interfaces within the area except those on which BFD is explicitly disabled. Given this rule, BFD is enabled on Gigabit Ethernet interfaces 1/0/0/0 and 2/0/0/0 and is disabled on interface 3/0/0/0.

In area 1, BFD is enabled on Gigabit Ethernet interface 5/0/0/0 only. Because BFD is not enabled at the area level nor explicitly enabled on interface 4/0/0/0, it is disabled on this interface.

```
router ospfv3 san_jose
    area 0
        bfd fast-detect
        ...
        int gige 1/0/0/0
        !
        int gige 2/0/0/0
            ...
int gige 3/0/0/0
            bfd fast-detect disable
!
    !
    area 1
        int gige 4/0/0/0
        !
        int gige 5/0/0/0
            bfd fast-detect
        !
    !
```

Related Commands

Command	Description
address-family ipv4 unicast (BFD), on page 627	Enables BFD fast-detection on a specific IPV4 unicast destination address prefix and on the forwarding next-hop address.
bfd minimum-interval, on page 655	Specifies the minimum control packet interval for BFD sessions for the corresponding BFD configuration scope.
bfd multiplier, on page 662	Sets the BFD multiplier.
show bfd, on page 687	Displays BFD information for a specific location.

bfd minimum-interval

To specify the minimum control packet interval for BFD sessions for the corresponding BFD configuration scope, use the **bfd minimum-interval** command in the appropriate configuration mode. To return the router to the default setting, use the **no** form of this command.

bfd minimum-interval *milliseconds*
no bfd minimum-interval [*milliseconds*]

Syntax Description

milliseconds	Interval between sending BFD hello packets to the neighbor. The range is 15 to 30000 milliseconds. For MPLS-TE, the range is 15 to 200 milliseconds. For GRE tunnel, the range is 150 to 30000 milliseconds. The default is 150 milliseconds.

Command Default

BGP *interval*: 50 milliseconds

IS-IS *interval*: 150 milliseconds

OSPF and OSPFv3 *interval*: 150 milliseconds

MPLS-TE *interval*: 15 milliseconds

PIM *interval*: 150 milliseconds

Command Modes

Router configuration

Interface configuration

MPLS TE configuration

Router configuration

Area configuration

Area interface configuration

Router configuration

Area configuration

Interface configuration

Tunnel configuration

Command History

Release	Modification
Release 3.7.2	This command was introduced.
Release 4.0.0	The **bfd minimum-interval** command was supported in the following modes: OSPFv3 router configuration, OSPFv3 area configuration, and OSPFv3 area interface configuration.
Release 6.5.1	The **bfd minimum-interval** command was supported in the tunnel configuration mode.

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.

In OSPF and OSPFv3 environments, the setting of the **bfd minimum-interval** command is inherited from the highest-level configuration mode in which the command was configured. From the lowest to the highest configuration modes, the inheritance rules are as follows:

- If you configure the minimum interval in area interface configuration mode, the updated interval affects the BFD sessions on the specified interface only.
- If you configure the minimum interval in area configuration mode, the updated interval affects the BFD sessions on all interfaces in the specified area.
- If you configure the minimum interval in router configuration mode, the updated interval affects the BFD sessions in all areas and all associated interfaces in the specified routing process.

If desired, you can override these inheritance rules by explicitly configuring the **bfd minimum-interval** command for a specific area interface or area.



Note When multiple applications share the same BFD session, the application with the most aggressive timer wins locally. Then, the result is negotiated with the peer router.

Keep the following router-specific rules in mind when configuring the minimum BFD interval:

- The maximum rate in packets-per-second (pps) for BFD sessions is linecard-dependent. If you have multiple linecards supporting BFD, then the maximum rate for BFD sessions per system is the supported linecard rate multiplied by the number of linecards.

  The maximum rate for BFD sessions per linecard is 9600 pps.
- If a session is running in asynchronous mode without echo, then PPS used for this session is (1000 / asynchronous interval in milliseconds).
- If a session is running in asynchronous mode with echo, then PPS used for this session is (1000 / echo interval in milliseconds).

  This is calculated as: 1000 / value of the **bfd minimum-interval** command.



Note The rate for BFD sessions on bundle member links is calculated differently. For more information, see the **bfd address-family ipv4 minimum-interval** command.

- The maximum number of all BFD sessions per linecard is 1024.
- The maximum number of all BFD sessions per linecard is 1440.
- When asynchronous mode is available, the minimum interval must be greater than or equal to 15 milliseconds for up to 100 sessions on the line card. If you are running the maximum of 1024 sessions, the failure detection interval must be greater than or equal to 150 milliseconds.
- When asynchronous mode is available, the minimum interval must be greater than or equal to 250 milliseconds, with a multiplier of 3 for up to 100 sessions per line card.

- When asynchronous mode is available, the minimum interval must be greater than or equal to 15 milliseconds for up to 100 sessions on the line card. If you are running the maximum of 1440 sessions, the failure detection interval must be greater than or equal to 150 milliseconds.
- When echo mode is available, the minimum interval must be greater than or equal to 15 milliseconds for up to 100 sessions on the line card. If you are running the maximum of 1024 sessions, the failure detection interval must be less than or equal to 150 milliseconds.
- When echo mode is available, the minimum interval must be 50 milliseconds with a multiplier of 3.
- When echo mode is available, the minimum interval must be greater than or equal to 15 milliseconds for up to 100 sessions on the line card. If you are running the maximum of 1440 sessions, the failure detection interval must be less than or equal to 150 milliseconds.

Task ID

Task ID	Operations
bgp	read, write
isis	read, write
mpls-te	read, write
ospf	read, write
multicast	read, write

Examples

The following example shows how to set the BFD minimum interval for a BGP routing process:

```
RP/0/RSP0/CPU0:router(config)# router bgp 6500
RP/0/RSP0/CPU0:router(config-bgp)# bfd minimum-interval 275
```

The following example shows the configuration of an OSPFv3 routing process named san_jose. The example shows two areas, each of which includes Gigabit Ethernet interfaces. In area 0, the minimum interval is set to 200 at the area level, which means that by virtue of the inheritance rules, the same value is set on all interfaces within the area except those on which a different value is explicitly configured. Given this rule, Gigabit Ethernet interface 1/0/0/0 uses the interval of 200, which is inherited from the area, while interface 2/0/0/0 uses the explicitly configured value of 300.

In area 1, the minimum interval is not configured at the area or interface levels, which means that interfaces 3/0/0/0 and 4/0/0/0 use the default interval of 150.

```
router ospfv3 san_jose
bfd fast-detect
    area 0
bfd minimum-interval 200
int gige 1/0/0/0
        !
int gige 2/0/0/0
```

```
bfd minimum-interval 300
        !
    !
    area 1
int gige 3/0/0/0
        !
int gige 4/0/0/0
!
    !
```

Related Commands

Command	Description
address-family ipv4 unicast (BFD), on page 627	Enables BFD fast-detection on a specific IPV4 unicast destination address prefix and on the forwarding next-hop address.
bfd fast-detect, on page 652	Enables BFD to detect failures in the path between adjacent forwarding engines.
bfd multiplier, on page 662	Sets the BFD multiplier.
show bfd, on page 687	Displays BFD information for a specific location.

bfd mode

To enable the option to use Cisco or IETF mode for BFD over bundle, use the **bfd mode** command in interface configuration mode. To disable the option to use Cisco or IETF mode for BFD over bundle, use the **no** form of this command.

bfd mode {**cisco** | **ietf**}

no bfd mode

Syntax Description

cisco	Specifies the use of Cisco mode for BFD over bundle.
ietf	Specifies the use of IETF mode for BFD over bundle.

Command Default

The default member mode is **cisco**.

Command Modes

Interface configuration (config-if)

Command History

Release	Modification
Release 5.3.1	This command was introduced.

Usage Guidelines

If the BFD mode is configured when the bundle is being created, the configuration goes through. This is because, both the BFD state as well as the bundle state are 'down' during bundle creation. To apply the mode change for existing sessions, bring down and then recreate the BFD sessions for that bundle.

A BFD session on the member interfaces can belong to only one mode (Cisco or IETF mode). Mix of the modes within the same bundle is not supported. This command is supported on bundle interfaces only.

Task ID

Task ID	Operations
bundle	read, write

Examples

This example shows how to enable **ietf** mode for the BFD session on an Ethernet bundle interface:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)#  interface bundle-ether 1
RP/0/RSP0/CPU0:router(config-if)#   bfd mode ietf
```

bfd multipath include location

To include specific linecards to host BFD multiple path sessions, use the **bfd multipath include location** command in the global configuration mode. To remove the configuration, use the **no** form of this command.

bfd multipath include location *node-id*
no bfd multipath include location *node-id*

Syntax Description

location *node-id*	Configures BFD multipath on the specified location. The *node-id* variable is mentioned in the *rack/slot/module* notation.

Command Default

No default behavior or values

Command Modes

Global configuration

Command History

Release	Modification
Release 4.2.0	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
ospf	read, write
isis	read, write
mpls-te	read, write

Example

This example shows how to run the **bfd multipath include location** command on a specific location:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# bfd multipath include location 0/5/CPU0
```

Related Commands

Command	Description
bfd, on page 629	Enters BFD configuration mode.
show bfd multipath, on page 700	Displays information regarding BFD multipath sessions.

bfd multiplier

To set the Bidirectional Forwarding Detection (BFD) multiplier, use the **bfd multiplier** command in the appropriate configuration mode. To return the router to the default setting, use the **no** form of this command.

bfd multiplier *multiplier*
no bfd multiplier [*multiplier*]

Syntax Description

multiplier Number of times a packet is missed before BFD declares the neighbor down. The ranges are as follows:

- BGP—2 to 16
- IS-IS—2 to 50
- MPLS-TE—2 to 10
- OSPF and OSPFv3—2 to 50
- PIM—2 to 50
- GRE tunnel—3 to 50

Command Default

The default multiplier is 3.

Command Modes

Router configuration

Interface configuration

MPLS-TE configuration

Router configuration

Area configuration

Area interface configuration

Router configuration

Area configuration

Area interface configuration

Interface configuration

Tunnel configuration

Command History

Release	Modification
Release 3.7.2	This command was introduced.
Release 4.0.0	The **bfd multiplier** command was supported in the following modes: OSPFv3 router configuration, OSPFv3 area configuration, and OSPFv3 area interface configuration.

Release	Modification
Release 6.5.1	The **bfd multiplier** command was supported in tunnel configuration mode.

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.

In OSPF and OSPFv3 environments, the setting of the **bfd multiplier** command is inherited from the highest-level configuration mode in which the command was configured. From the lowest to the highest configuration modes, the inheritance rules are as follows:

- If you configure a multiplier in area interface configuration mode, the updated multiplier affects the BFD sessions on the specified interface only.
- If you configure a multiplier in area configuration mode, the updated multiplier affects the BFD sessions on all interfaces in the specified area.
- If you configure a multiplier in router configuration mode, the updated multiplier affects the BFD sessions in all areas and all associated interfaces in the specified routing process.

If desired, you can override these inheritance rules by explicitly configuring the **bfd multiplier** command for a specific area interface or area.

If the multiplier is changed using the **bfd multiplier** command, the new value is used to update all existing BFD sessions for the protocol (BGP, IS-IS, MPLS-TE, OSPF, or OSPFv3) .

Task ID

Task ID	Operations
bgp	read, write
isis	read, write
mpls-te	read, write
ospf	read, write
multicast	read, write

Examples

The following example shows how to set the BFD multiplier in a BGP routing process:

```
RP/0/RSP0/CPU0:router(config)# router bgp 65000
RP/0/RSP0/CPU0:router(config-bgp)# bfd multiplier 2
```

The following example shows the configuration of an OSPFv3 routing process named san_jose. The example shows two areas, each of which includes Gigabit Ethernet interfaces. In area 0, the multiplier is set to 5 at the area level, which means that by virtue of the inheritance rules, the same value is set

on all interfaces within the area except those on which a different value is explicitly configured. Given this rule, Gigabit Ethernet interface 1/0/0/0 uses a multiplier of 5, which is inherited from the area, while interface 2/0/0/0 uses the explicitly configured multiplier of 2.

In area 1, a multiplier is not configured at the area or interface levels, which means that interfaces 3/0/0/0 and 4/0/0/0 use the default value of 3.

```
router ospfv3 san_jose
bfd fast-detect
    area 0
bfd multiplier 5
int gige 1/0/0/0
        !
int gige 2/0/0/0
bfd multiplier 2
        !
    !
    area 1
int gige 3/0/0/0
        !
int gige 4/0/0/0
!
    !
```

Related Commands

Command	Description
address-family ipv4 unicast (BFD), on page 627	Enables BFD fast-detection on a specific IPV4 unicast destination address prefix and on the forwarding next-hop address.
bfd fast-detect, on page 652	Enables BFD to detect failures in the path between adjacent forwarding engines.
bfd minimum-interval, on page 655	Specifies the minimum control packet interval for BFD sessions for the corresponding BFD configuration scope.
show bfd, on page 687	Displays BFD information for a specific location.

bfd bundle per-member mode

To enable the option to use Cisco or IETF mode for BFD over per-bundle member link, use the **bfd bundle per-member mode** command in router configuration mode. To disable the option to use Cisco or IETF mode for BFD over per-bundle member link, use the **no** form of this command.

bfd bundle per-member mode {**cisco** | **ietf**}

no bfd bundle per-member mode {**cisco** | **ietf**}

Syntax Description

cisco	Specifies the use of Cisco mode for BFD over per-bundle member link.
ietf	Specifies the use of IETF mode for BFD over per-bundle member link.

Command Default

The default member mode is **cisco**.

Command Modes

Router configuration

Command History

Release	Modification
Release 5.2.0	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.

The mode change is applied only for new sessions. To apply mode change for existing sessions, delete and then recreate the sessions.

A BFD session on the member interfaces can belong to only one mode (Cisco or IETF mode). Mix of the modes within the same bundle is not supported.

Task ID

Task ID	Operations
bfd	read, write

Examples

This example shows how to enable **cisco** mode for BFD over per-bundle member link:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# bfd bundle per-member mode cisco
```

clear bfd counters

To clear Bidirectional Forwarding Detection (BFD) counters, use the **clear bfd counters** command in EXEC mode.

clear bfd counters {**ipv4** | [**singlehop** | **multihop**] | **ipv6** | [**singlehop** | **multihop**] | **all** | **label**} [**packet**] [**timing**] [**interface** *type interface-path-id*] **location** *node-id*

Syntax Description

ipv4	(Optional) Clears BFD over IPv4 information only.
ipv6	(Optional) Clears BFD over IPv6 information only.
singlehop	(Optional) Clears BFD singlehop information only.
multihop	(Optional) Clears BFD multihop information only.
all	(Optional) Clears both BFD over IPv4 and BFD over IPv6 information.
packet	(Optional) Specifies that packet counters are cleared.
timing	(Optional) Specifies that timing counters are cleared.
interface	(Optional) Specifies the interface from which the BFD packet counters are cleared.
type	Specifies the 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.
location *node-id*	Clears BFD counters from the specified location. The *node-id* argument is entered in the *rack/slot/module* notation.

Command Default

The default is the default address family identifier (AFI) that is set by the **set default-afi** command, IPv4 or IPv6.

Command Modes

EXEC

Command History

Release	Modification
Release 3.7.2	This command was introduced.
Release 4.0.0	Support for the **ipv6** keyword was added.
Release 4.2.0	Support for the singlehop and multihop keywords were added.

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.

For the *interface-path-id* argument, use the following guidelines:

- If specifying a physical interface, the naming notation is *rack*/*slot*/*module*/*port*. The slash between values is required as part of the notation. An explanation of each component of the naming notation is as follows:
  - *rack*: Chassis number of the rack.
  - *slot*: Physical slot number of the line card.
  - *module*: Module number. A physical layer interface module (PLIM) is always 0.
  - *port*: Physical port number of the interface.
- If specifying a virtual interface, the number range varies, depending on interface type.

Both IPv4 and IPv6 BFD sessions can run simultaneously on the same line card.

Examples

The following example shows how to clear the BFD IPv6 packet counters on a POS interface:

```
RP/0/RSP0/CPU0:router# clear bfd counters packet ipv6 interface POS 0/1/0/0 location 0/1/cpu0
```

The following example shows how to clear the BFD IPv4 timing counters:

```
RP/0/RSP0/CPU0:router# clear bfd counters ipv4 timing location 0/5/cpu0
```

Related Commands

Command	Description
bfd fast-detect, on page 652	Enables BFD to detect failures in the path between adjacent forwarding engines.
bfd minimum-interval, on page 655	Specifies the minimum control packet interval for BFD sessions for the corresponding BFD configuration scope.
bfd multiplier, on page 662	Sets the BFD multiplier.
show bfd, on page 687	Displays BFD information for a specific location.

clear bfd dampening

To reset BFD dampening counters, use the **clear bfd dampening** command in EXEC mode.

clear bfd dampening {[**interface** *type interface-path-id*] | **location** *node-id all* | **type** [**bundle-logical** | **bundle-per-member**]}

Syntax Description

interface	Specifies the interface from which the BFD dampening sessions are cleared.
type	Specifies the 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.
location *node-id*	Clears BFD dampening sessions from the specified location. The *node-id* argument is entered in the *rack/slot/module* notation.
all	(Optional) Clears BFD dampening session on all locations.
type	Specifies the BFD session type. For more information, use the question mark (**?**) online help function.
bundle-per-member	Clears BFD dampening over the member links of BFD over bundle sessions.
bundle-logical	Clears BFD dampening of the BFD over logical bundle session.

Command Modes

EXEC

Command History

Release	Modification
Release 4.3.2	This command was introduced.
Release 5.3.0	Added **type** and **interface** keywords.

Usage Guidelines

For the *interface-path-id* argument, use the following guidelines:

- If specifying a physical interface, the naming notation is *rack*/*slot*/*module*/*port*. The slash between values is required as part of the notation. An explanation of each component of the naming notation is as follows:

- *rack*: Chassis number of the rack.
- *slot*: Physical slot number of the line card.
- *module*: Module number. A physical layer interface module (PLIM) is always 0.
- *port*: Physical port number of the interface.

- If specifying a virtual interface, the number range varies, depending on interface type.

Examples

The following example shows how to clear BFD dampening:

- To clear BFD dampening on all BFD sessions:

```
RP/0/RSP0/CPU0:router# clear bfd dampening location all
```

- To clear BFD dampening on BFD sessions over bundle interfaces:

```
RP/0/RSP0/CPU0:router# clear bfd dampening type bundle-per-member
```

- To clear bfd dampening on BFD sessions over bundle logical interfaces:

```
RP/0/RSP0/CPU0:router# clear bfd dampening type bundle-logical
```

- To clear bfd dampening on BFD sessions over specific interface:

```
RP/0/RSP0/CPU0:router# clear bfd dampening interface gigabitEthernet 0/3/0/0
```

clear bfd dampening log sample

The following is a sample syslog output of the **clear bfd dampening** command:

```
LC/0/3/CPU0:Oct  6 16:52:40.632 : bfd_agent[127]: %L2-BFD-6-SESSION_DAMPENING_CLR : Dampening
 for BFD session to neighbor 10.1.1.2
on interface GigabitEthernet0/3/0/0 has been cleared
```

echo disable

To disable echo mode on a router or on an individual interface or bundle, use the **echo disable** command in Bidirectional Forwarding Detection (BFD) configuration mode. To return the router to the default configuration where echo mode is enabled, use the **no** form of this command.

echo disable
no echo disable

Syntax Description

This command has no keywords or arguments.

Command Default

No default behavior or values

Command Modes

BFD configuration

BFD 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

If you are using BFD with Unicast Reverse Path Forwarding (uRPF), you need to use the **echo disable** command to disable echo mode; otherwise, echo packets are rejected.



Note To enable or disable IPv4 uRPF checking on an IPv4 interface, use the **[no] ipv4 verify unicast source reachable-via** command in interface configuration mode. To enable or disable loose IPv6 uRPF checking on an IPv6 interface, use the **[no] ipv6 verify unicast source reachable-via any** command in interface configuration mode.

Task ID

Task ID	Operations
bgp	read, write
ospf	read, write
isis	read, write
mpls-te	read, write

Examples

The following example shows how to disable echo mode on a router:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# bfd
RP/0/RSP0/CPU0:router(config-bfd)# echo disable
```

The following example shows how to disable echo mode on an individual interface:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# bfd
RP/0/RSP0/CPU0:router(config-bfd)# interface gigabitethernet 0/1/0/0
RP/0/RSP0/CPU0:router(config-bfd-if)# echo disable
```

Related Commands

Command	Description
bfd, on page 629	Enters BFD configuration mode.
interface (BFD), on page 680	Enters BFD interface configuration mode.
ipv4 verify unicast source reachable-via	Enables and disables IPv4 uRPF checking on an IPv4 interface.
ipv6 verify unicast source reachable-via any	Enables and disables loose IPv6 uRPF checking on an IPv6 interface.
show bfd, on page 687	Displays BFD information.

echo ipv4 source

To specify the IP address that you want to use as the source address for BFD echo packets, use the **echo ipv4 source** command in BFD or BFD interface configuration mode. To return to the default, use the **no** form of this command.

echo ipv4 source *ip-address*
no echo ipv4 source *ip-address*

Syntax Description

ip-address	32-bit IPv4 address in dotted-decimal format (A.B.C.D).

Command Default

The IP address of the output interface, or the IP address in the **router-id** command (if configured), is the default address used for an echo packet when the **echo ipv4 source** command is not configured.

Command Modes

BFD configuration

BFD interface configuration

Command History

Release	Modification
Release 3.9.0	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.

If you do not configure the IPv4 source address for echo packets, then BFD uses the IP address of the output interface or the address in the **router-id** command if specified.

You can override the default address for BFD echo packets by specifying an IPv4 source address for echo packets globally for all BFD sessions on the router and at an individual interface. Specifying the IP address at an individual interface will override any value specified globally for BFD on the router.

Task ID

Task ID	Operations
bgp	read, write
ospf	read, write
isis	read, write
mpls-te	read, write

Examples

The following example shows how to specify the IP address 10.10.10.1 as the source address for BFD echo packets for all BFD sessions on the router:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# bfd
RP/0/RSP0/CPU0:router(config-bfd)# echo ipv4 source 10.10.10.1
```

The following example shows how to specify the IP address 10.10.10.1 as the source address for BFD echo packets on an individual Gigabit Ethernet interface:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# bfd
RP/0/RSP0/CPU0:router(config-bfd)# interface gigabitethernet 0/1/0/0
RP/0/RSP0/CPU0:router(config-bfd-if)# echo ipv4 source 10.10.10.1
```

The following example shows how to specify the IP address 10.10.10.1 as the source address for BFD echo packets on an individual Packet-over-SONET (POS) interface:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# bfd
RP/0/RSP0/CPU0:router(config-bfd)# interface pos 0/1/0/0
RP/0/RSP0/CPU0:router(config-bfd-if)# echo ipv4 source 10.10.10.1
```

Related Commands

Command	Description
bfd, on page 629	Enters BFD configuration mode.
echo disable, on page 670	Disables echo mode on a router or on an individual interface or bundle.

echo latency detect

To enable latency detection for BFD echo packets, use the **echo latency detect** command in BFD configuration mode. To return to the default, use the **no** form of this command.

echo latency detect [**percentage** *percent-value* [**count** *packet-count*]]
no echo latency detect [**percentage** *percent-value* [**count** *packet-count*]]

Syntax Description

percentage *percent-value*	(Optional) Percentage of the echo failure detection time to be detected as bad latency. The range is 100 to 250. The default is 100.
count *packet-count*	(Optional) Number of consecutive packets received with the detected bad latency that will take down a BFD session. The range is 1 to 10. The default is 1.

Command Default

Echo latency detection is disabled.

Command Modes

BFD configuration

Command History

Release	Modification
Release 4.0.0	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.

Note Latency detection is only valid where echo mode is supported for BFD. However, it is not supported on bundle interfaces.

Without latency detection, standard BFD echo failure detection tracks only the absence of receipt of echo packets within a period of time based on a counter. However, this standard echo failure detection does not address latency between transmission and receipt of any specific echo packet, which can build beyond desired tolerances over the course of the BFD session.

When latency detection is enabled, a percentage is multiplied to the echo failure detection value (I x M x %), and the roundtrip delay is computed for the echo packet. If this delay is greater than (I x M x %), then the BFD session is taken down.

If you have specified a packet count, then the system tracks the number of packets received back-to-back with bad latency before taking down the session.

Task ID

Task ID	Operations
bgp	read, write

Task ID	Operations
ospf	read, write
isis	read, write
mpls-te	read, write

Examples

In the following examples, consider that the BFD minimum interval is 50 ms, and the multiplier is 3 for the BFD session.

The following example shows how to enable echo latency detection using the default values of 100% of the echo failure period (I x M) for a packet count of 1. In this example, when one echo packet is detected with a roundtrip delay greater than 150 ms, the session is taken down:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# bfd
RP/0/RSP0/CPU0:router(config-bfd)# echo latency detect
```

The following example shows how to enable echo latency detection based on 200% (two times) of the echo failure period for a packet count of 1. In this example, when one packet is detected with a roundtrip delay greater than 300 ms, the session is taken down:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# bfd
RP/0/RSP0/CPU0:router(config-bfd)# echo latency detect percentage 200
```

The following example shows how to enable echo latency detection based on 100% of the echo failure period for a packet count of 3. In this example, when three consecutive echo packets are detected with a roundtrip delay greater than 150 ms, the session is taken down:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# bfd
RP/0/RSP0/CPU0:router(config-bfd)# echo latency detect percentage 100 count 3
```

Related Commands

Command	Description
bfd, on page 629	Enters BFD configuration mode.
bfd minimum-interval, on page 655	Specifies the minimum control packet interval for BFD sessions for the corresponding BFD configuration scope.
bfd multiplier, on page 662	Sets the BFD multiplier.
echo startup validate, on page 676	Enables verification of the echo packet path before starting a BFD session.

echo startup validate

To enable verification of the echo packet path before starting a BFD session, use the **echo startup validate** command in BFD configuration mode. To return to the default, use the **no** form of this command.

echo startup [**force**]
no echo startup [**force**]

Syntax Description

force	(Optional) Ignores the remote 'Required Min Echo RX Interval' setting.

Command Default

Echo startup validation is disabled.

Command Modes

BFD configuration

Command History

Release	Modification
Release 4.0.0	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.

Note Echo validation is only valid where echo mode is supported for BFD. However, it is not supported on bundle interfaces.

When a BFD session is down and the **echo startup validate** command is configured, an echo packet is periodically transmitted on the link while it is down to verify successful transmission within the configured latency before allowing the BFD session to change state.

Without the **force** option, the echo validation test only runs if the last received control packet contains a non-zero “Required Min Echo RX Interval” value. When the **force** keyword is configured, the echo validation test runs regardless of this value.

Task ID

Task ID	Operations
bgp	read, write
ospf	read, write
isis	read, write
mpls-te	read, write

Examples

The following example shows how to enable echo startup validation for BFD sessions on non-bundle interfaces if the last received control packet contains a non-zero “Required Min Echo RX Interval” value:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# bfd
RP/0/RSP0/CPU0:router(config-bfd)# echo startup validate
```

The following example shows how to enable echo startup validation for BFD sessions on non-bundle interfaces regardless of the “Required Min Echo RX Interval” value in the last control packet:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# bfd
RP/0/RSP0/CPU0:router(config-bfd)# echo startup validate force
```

Related Commands

Command	Description
bfd, on page 629	Enters BFD configuration mode.
echo latency detect, on page 674	Enables latency detection for BFD echo packets.

hw-module bfd-hw-offload

To enable BFD hardware offload mode, use the **hw-module bfd-hw-offload** command in admin mode. To disable BFD hardware offload mode, use the **no** form of this command.

Note
- The BFD hardware offload mode is supported on Cisco ASR 9000 Enhanced Ethernet Line Card and Cisco ASR 9000 High Density 100GE Ethernet Line Card.

hw-module bfd-hw-offload enable location *line card location*
no hw-module bfd-hw-offload enable location *line card location*

Syntax Description

enable	Enables BFD hardware offload mode. Only BFD IPv4 and IPv6 sessions over physical and VLAN interfaces (in asynchronous mode) and BFD over MPLS-TP LSP single path sessions are supported.
line card location	Specify the line card location. Only the ASR 9000 Enhanced Ethernet Line Card supports BFD hardware offload mode.

Command Default

BFD hardware offload mode is disabled.

Command Modes

Global configuration.

Command History

Release	Modification
Release 5.1.2	This command was introduced.

Usage Guidelines

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

- BFD hardware offload mode is disabled by default. You can enable it in admin mode.
- You must reload the line cards after enabling or disabling BFD hardware offload mode for the configuration to take effect.
- BFD hardware offload is not supported on virtual interfaces.

Task ID

Task ID	Operations
root	read, write

Examples

The below example shows how to enable BFD hardware offload mode on the ASR 9000 Enhanced Ethernet Line Card in the location 0/0/CPU0:

```
RP/0/RSP0/CPU0:router# admin
RP/0/RSP0/CPU0:router(admin)# hw-module bfd-hw-offload enable location 0/0/CPU0
```

interface (BFD)

To enter Bidirectional Forwarding Detection (BFD) interface configuration mode, where you can disable echo mode on an interface, use the **interface** command in BFD configuration mode. To return to BFD configuration mode, use the **no** form of this command.

interface *type interface-path-id*
no interface *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

No default behavior or values

Command Modes

BFD 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

For the *interface-path-id* argument, use the following guidelines:

- If specifying a physical interface, the naming notation is *rack*/*slot*/*module*/*port*. The slash between values is required as part of the notation. An explanation of each component of the naming notation is as follows:
  - *rack*: Chassis number of the rack.
  - *slot*: Physical slot number of the line card.
  - *module*: Module number. A physical layer interface module (PLIM) is always 0.
  - *port*: Physical port number of the interface.
- If specifying a virtual interface, the number range varies, depending on interface type.

If you are using BFD with Unicast Reverse Path Forwarding (uRPF) on a particular interface, then you need to use the **echo disable** command in BFD interface configuration mode to disable echo mode on that interface; otherwise, echo packets are rejected by the interface.

Note To enable or disable IPv4 uRPF checking on an IPv4 interface, use the **[no] ipv4 verify unicast source reachable-via** command in interface configuration mode. To enable or disable loose IPv6 uRPF checking on an IPv6 interface, use the **[no] ipv6 verify unicast source reachable-via any** command in interface configuration mode.

Task ID

Task ID	Operations
bgp	read, write
ospf	read, write
isis	read, write
mpls-te	read, write

Examples

The following example shows how to enter BFD interface configuration mode for a Gigabit Ethernet interface:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# bfd
RP/0/RSP0/CPU0:router(config-bfd)# interface gigabitethernet 0/1/0/0
RP/0/RSP0/CPU0:router(config-bfd-if)#
```

The following example shows how to enter BFD interface configuration mode for a Packet-over-SONET/SDH (POS) interface:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# bfd
RP/0/RSP0/CPU0:router(config-bfd)# interface pos 0/1/0/0
RP/0/RSP0/CPU0:router(config-bfd-if)#
```

Related Commands

Command	Description
bfd, on page 629	Enters BFD configuration mode.
echo disable, on page 670	Disables echo mode on an individual interface or on the entire router.
ipv4 verify unicast source reachable-via	Enables and disables IPv4 uRPF checking on an IPv4 interface.

Command	Description
ipv6 verify unicast source reachable-via any	Enables and disables loose IPv6 uRPF checking on an IPv6 interface.
show bfd, on page 687	Displays BFD information.

ipv6 checksum

To enable and disable IPv6 checksum calculations on BFD UDP packets globally or on a BFD interface, use the **ipv6 checksum** command in Bidirectional Forwarding Detection (BFD) or BFD interface configuration mode. To return to the default, use the **no** form of this command.

BFD Configuration
ipv6 checksum disable
no ipv6 checksum disable

BFD Interface Configuration
ipv6 checksum [**disable**]
no ipv6 checksum [**disable**]

Syntax Description

disable	(Optional for BFD interface configuration only) Disables IPv6 checksum calculations.

Command Default

IPv6 checksum calculations on BFD UDP packets is disabled.

Command Modes

BFD configuration

BFD interface configuration

Command History

Release	Modification
Release 4.0.0	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.

IPv6 checksum calculations for UDP packets are disable by default for BFD sessions. You can enable IPv6 checksum support either globally for all BFD sessions, or on an individual interface.

Task ID

Task ID	Operations
bgp	read, write
ospf	read, write
isis	read, write
mpls-te	read, write

Examples

The following example shows how to enable IPv6 checksum calculations for UDP packets for all BFD sessions on the router:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# bfd
RP/0/RSP0/CPU0:router(config-bfd)# no ipv6 checksum disable
```

The following example shows how to disable IPv6 checksum calculations for UDP packets for all BFD sessions on the router:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# bfd
RP/0/RSP0/CPU0:router(config-bfd)# ipv6 checksum disable
```

The following example shows how to enable echo mode for BFD sessions on an individual interface:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# bfd
RP/0/RSP0/CPU0:router(config-bfd)# interface gigabitethernet 0/1/0/0
RP/0/RSP0/CPU0:router(config-bfd-if)# ipv6 checksum
```

The following example shows how to disable echo mode for BFD sessions on an individual interface:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# bfd
RP/0/RSP0/CPU0:router(config-bfd)# interface gigabitethernet 0/1/0/0
RP/0/RSP0/CPU0:router(config-bfd-if)# ipv6 checksum disable
```

Related Commands

Command	Description
bfd, on page 629	Enters BFD configuration mode.
interface (BFD), on page 680	Enters BFD interface configuration mode, where you can disable echo mode on an interface.
show bfd, on page 687	Displays BFD information for a specific location.

multihop ttl-drop-threshold

To specify the maximum time to live (TTL) value for multihop sessions per system, use the **multihop ttl-drop-threshold** command in the BFD configuration mode. To return to the default, use the **no** form of this command.

multihop ttl-drop-threshold *value*
no multihop ttl-drop-threshold *value*

Syntax Description

value	Specifies the configurable range of values for TTL. It ranges from 0 to 254.

Command Default

No default behavior or values

Command Modes

BFD configuration

Command History

Release	Modification
Release 4.2.0	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.

When the TTL of a BFD packet is received on the router, the value is decremented by one and then compared with the configured ttl-drop-threshold.

- If the TTL of a BFD packet after decrementing is less than the configured ttl-drop-threshold, the packet will be dropped.
- If the TTL of a BFD packet received on the router after decrementing is greater than or equal to the configured ttl-drop-threshold, the packet will be processed.



Note This configuration command is only applicable for BFD multihop sessions.

Task ID

Task ID	Operation
bgp	read, write
ospf	read, write
isis	read, write

Task ID	Operation
mpls-te	read, write

Example

This example shows how to set the maximum TTL value as 2 using the **multihop ttl-drop-threshold** command:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# bfd multihop ttl-drop-threshold 2
```

Related Commands

Command	Description
show bfd counters, on page 692	Displays BFD counter information.
show bfd multipath, on page 700	Displays information regarding BFD multipath sessions.

show bfd

To display Bidirectional Forwarding Detection (BFD) information for a specific location, use the **show bfd** command in EXEC mode.

show bfd [**ipv4** | [**singlehop** | **multihop**] | **ipv6** [**singlehop** | **multihop**]| **all**|**label**]**interface**[**destination** | **source**] [**location** *node-id*]

Syntax Description

ipv4	(Optional) Displays BFD over IPv4 information only.
multihop	(Optional) Displays BFD multihop information only.
singlehop	(Optional) Displays BFD singlehop information only.
ipv6	(Optional) Displays BFD over IPv6 information only.
all	(Optional) Displays both BFD over IPv4 and BFD over IPv6 information.
label	(Optional) Displays the BFD label information.
interface	Specifies the BFD interface.
destination	(Optional) Specifies the destination IPv4 unicast address.
source	(Optional) Specifies the source IPv4 unicast address.
location *node-id*	Displays BFD information for the specified location. The *node-id* argument is entered in the *rack/slot/module* notation.

Command Default

The default is the default address family identifier (AFI) that is set by the **set default-afi** command, IPv4 or IPv6.

Command Modes

EXEC

Command History

Release	Modification
Release 3.7.2	This command was introduced.
Release 4.0.0	Support for the **ipv6** keyword was added.
Release 4.2.0	Support for **multihop** keyword was added.

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

Task ID	Operations
ospf	read
isis	read
mpls-te	read

Examples

The following example shows the output from the **show bfd** command:

```
RP/0/RSP0/CPU0:router# show bfd

IPV4 Sessions Up: 0, Down: 0, Total: 0
```

The following example shows the output from the **show bfd all** command:

```
RP/0/RSP0/CPU0:router# show bfd all

IPv4:
-----
IPV4 Sessions Up: 20, Down: 0, Unknown/Retry: 2, Total: 22
IPv6:
-----
IPV6 Sessions Up: 128, Down: 2, Unknown/Retry: 1, Total: 131
Label:
-----
Label Sessions Up: 10, Down: 0, Unknown/Retry: 1, Total: 11
```

The following example shows the output from the **show bfd ipv4** command:

```
RP/0/RSP0/CPU0:router# show bfd ipv4

IPV4 Sessions Up: 0, Down: 0, Total: 0
```

The following example shows the output from the **show bfd ipv6** command:

```
RP/0/RSP0/CPU0:router# show bfd ipv6

IPV6 Sessions Up: 0, Down: 0, Total: 0
```

The following example shows the output from the **show bfd ipv4 location** command:

```
RP/0/RSP0/CPU0:router# show bfd ipv6 session detail interface tenGigE 0/0/0/0.100

I/f: TenGigE0/0/0/0.100, Location: 0/0/CPU0
Dest: 1001:1:1:1:1:1:1:2
Src: 1001:1:1:1:1:1:1:1
 State: UP for 0d:1h:37m:3s, number of times UP: 1
 Session type: PR/V6/SH
Received parameters:
 Version: 1, desired tx interval: 300 ms, required rx interval: 300 ms
 Required echo rx interval: 0 us, multiplier: 3, diag: None
```

```
 My discr: 2148076695, your discr: 2148075550, state UP, D/F/P/C/A: 0/0/0/1/0
Transmitted parameters:
 Version: 1, desired tx interval: 300 ms, required rx interval: 300 ms
 Required echo rx interval: 0 us, multiplier: 3, diag: None
 My discr: 2148075550, your discr: 2148076695, state UP, D/F/P/C/A: 0/1/0/1/0
Timer Values:
 Local negotiated async tx interval: 300 ms
 Remote negotiated async tx interval: 300 ms
 Desired echo tx interval: 0 s, local negotiated echo tx interval: 0 us
 Echo detection time: 0 us(0 us*3), async detection time: 900 ms(300 ms*3)
Local Stats:
 Intervals between async packets:
   Tx: Number of intervals=3, min=10 ms, max=2290 ms, avg=824 ms
       Last packet transmitted 5823 s ago
   Rx: Number of intervals=6, min=3 ms, max=2 s, avg=461 ms
       Last packet received 5823 s ago
Intervals between echo packets:
   Tx: Number of intervals=0, min=0 s, max=0 s, avg=0 s
       Last packet transmitted 0 s ago
   Rx: Number of intervals=0, min=0 s, max=0 s, avg=0 s
       Last packet received 0 s ago
 Latency of echo packets (time between tx and rx):
   Number of packets: 0, min=0 us, max=0 us, avg=0 us
Session owner information:
                          Desired               Adjusted
  Client                Interval   Multiplier Interval   Multiplier
  --------------------- ---------- ---------- ---------- ----------
  ipv6_static           300 ms     3          300 ms     3
  ipv6_static           300 ms     3          300 ms     3

H/W Offload Info:
 H/W Offload capability : Y, Hosted NPU     : 0/0/CPU0/NPU0
 Async Offloaded        : Y, Echo Offloaded : N
 Async rx/tx            : 7/4

Platform Info:
NPU ID: 0
Async RTC ID        : 1          Echo RTC ID        : 0
Async Feature Mask  : 0x8        Echo Feature Mask  : 0x0
Async Session ID    : 0x3c07     Echo Session ID    : 0x0
Async Tx Key        : 0x3c070801  Echo Tx Key        : 0x0
Async Tx Stats addr : 0x3f69e800   Echo Tx Stats addr : 0x0
Async Rx Stats addr : 0x4069e800   Echo Rx Stats addr : 0x0
```

Related Commands

Command	Description
bfd fast-detect, on page 652	Enables BFD to detect failures in the path between adjacent forwarding engines.
bfd minimum-interval, on page 655	Specifies the minimum control packet interval for BFD sessions for the corresponding BFD configuration scope.
bfd multiplier, on page 662	Sets the BFD multiplier.

show bfd client

To display Bidirectional Forwarding Detection (BFD) client information, use the **show bfd client** command in EXEC mode.

show bfd client [**detail**]

Syntax Description

detail	(Optional) Specifies detailed client information including number of sessions and client reconnects.

Command Default

Enter the **show bfd client** command without specifying the **detail** keyword to display summarized BFD client information.

Command Modes

EXEC

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 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
ospf	read
isis	read
mpls-te	read

Examples

The following example shows the output from the **show bfd client** command:

```
RP/0/RSP0/CPU0:router# show bfd client

Name            Node       Num sessions
--------------- ---------- --------------
bgp             0/RSP0/CPU0 0
isis            0/RSP0/CPU0 0
isis            0/RSP0/CPU0 0
```

Table 46: show bfd client Field Descriptions

Field	Description
Name	Name of the BFD client.

Field	Description
Node	Location of the BFD client.
Num sessions	Number of active sessions for the BFD client.

Related Commands

Command	Description
bfd fast-detect, on page 652	Enables BFD to detect failures in the path between adjacent forwarding engines.
bfd minimum-interval, on page 655	Specifies the minimum control packet interval for BFD sessions for the corresponding BFD configuration scope.
bfd multiplier, on page 662	Sets the BFD multiplier.
show bfd, on page 687	Displays BFD information for a specific location.

show bfd counters

To display Bidirectional Forwarding Detection (BFD) counter information, use the **show bfd counters** command in EXEC mode.

show bfd counters [**ipv4** | [**singlehop** | **multihop**] | **ipv6** [**singlehop** | **multihop**]| **all**| **label**] **packet** [**interface** *type interface-path-id*] **location** *node-id*

Syntax Description

ipv4	(Optional) Displays BFD over IPv4 information only.
ipv6	(Optional) Displays BFD over IPv6 information only.
singlehop	(Optional) Displays BFD singlehop information only.
multihop	(Optional) Displays BFD multihop information only.
all	(Optional) Displays both BFD over IPv4 and BFD over IPv6 information.
packet	Specifies that packet counters are displayed.
interface	(Optional) Specifies the interface for which to show counters.
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.
location *node-id*	Displays BFD counters from the specified location. The *node-id* argument is entered in the *rack/slot/module* notation.

Command Default

The default is the default address family identifier (AFI) that is set by the **set default-afi** command, IPv4 or IPv6.

Command Modes

EXEC

Command History

Release	Modification
Release 3.7.2	This command was introduced.
Release 4.0.0	Support for the **ipv6** keyword was added.
Release 4.2.0	Support for the **singlehop** and **multihop** keywords were added.

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.

For the *interface-path-id* argument, use the following guidelines:

- If specifying a physical interface, the naming notation is *rack*/*slot*/*module*/*port*. The slash between values is required as part of the notation. An explanation of each component of the naming notation is as follows:
  - *rack*: Chassis number of the rack.
  - *slot*: Physical slot number of the line card.
  - *module*: Module number. A physical layer interface module (PLIM) is always 0.
  - *port*: Physical port number of the interface.
- If specifying a virtual interface, the number range varies, depending on interface type.

Task ID

Task ID	Operations
bgp	read
ospf	read
isis	read
mpls-te	read

Examples

The following example shows the output from the **show bfd counters packet** command for both IPv4 and IPv6:

```
RP/0/RSP0/CPU0:router# show bfd counters packet all interface POS 0/1/0/0 location
 0/1/cpu0

Mon Nov  5 08:49:51.950 UTC
IPv4:
-----
 POS 0/1/0/0             Recv       Xmit                Recv       Xmit
      Async:             520        515       Echo:     9400       9400

IPv6:
-----
 POS 0/1/0/0             Recv       Xmit                Recv       Xmit
      Async:             237        237       Echo:        0          0
```

The following example shows the output from the **show bfd counters packet** command for IPv4:

```
RP/0/RSP0/CPU0:router# show bfd counters ipv4 packet

IPv4 Singlehop:
 GigabitEthernet0/0/1/2                  Recv       Xmit                 Recv       Xmit
      Async:                             4148       4137       Echo: (   47136)     80192
 GigabitEthernet0/1/1/2                  Recv       Xmit                 Recv       Xmit
      Async:                             116876     125756       Echo: ( 2268192)   2301312
```

```
 Bundle-Ether10                          Recv       Xmit                 Recv       Xmit
      Async:                             2          0       Echo:        0          0
 Bundle-Ether20                          Recv       Xmit                 Recv       Xmit
      Async:                             91         0       Echo:        0          0

IPv4 Multihop: (Src IP/Dst IP/Vrf Id)
 33.15.151.4/33.16.151.4/0x12345678      Recv       Xmit
      Async:                             0       570337
```

Table 47: show bfd counters packet Field Descriptions

Field	Description
Async	Number of asynchronous mode (control) packets that were received or transmitted on the specified interface.
Echo	Number of echo packets that were received or transmitted on the specified interface.

Related Commands

Command	Description
bfd fast-detect, on page 652	Enables BFD to detect failures in the path between adjacent forwarding engines.
bfd minimum-interval, on page 655	Specifies the minimum control packet interval for BFD sessions for the corresponding BFD configuration scope.
bfd multiplier, on page 662	Sets the BFD multiplier.

show bfd hw-offload

To display BFD hardware offload information, use the **show bfd hw-offload** command in EXEC mode.

show bfd hw-offload {**state location** *location-id* | **summary location** *location-id*}

Syntax Description

state	Displays if BFD hardware offload is enabled or disabled on the line card.
summary	Displays total number of sessions configured for each timer interval on the line card and network processor unit.
location-id	Specifies location-ID number of the line card.

Command Default

None.

Command Modes

EXEC

Command History

Release	Modification
Release 5.1.2	This command was introduced.

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes 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
root	read

This is the sample output from the **show bfd hw-offload** command when **state** is used:

```
RP/0/RSP0/CPU0:router# show bfd hw-offload state location 0/0/CPU0

BFD HW OFFLOAD Feature state:

0/0/CPU0
-----------------
Configuration State:   Enabled
Operational   State:   Enabled
```

This is the sample output from the **show bfd hw-offload** command when **summary** is used:

```
RP/0/RSP0/CPU0:router# show bfd hw-offload summary location 0/2/CPU0
BFD HW OFFLOAD Feature Summary:

0/2/CPU0
==========

The below available numbers per timer interval indicates the
```

```
max. sessions that can be configurd at that interval without
configuring any other session at any other interval.

After configuring, execute this CLI to get the remaining
available numbers.

              3.3ms   15ms    50ms    300ms   1s      2s      30s
-------------------------------------------------------------------------
Max LC Supp   600     2000    8000    8000    8000    8000    8000
Max NP Supp   300     1000    3000    3000    3000    3000    3000

-------------------------------------------------------------------------
LC:
----
Tx Used       0       0       8000    0       0       0       0
Rx Used       0       0       8000    0       0       0       0
Tx Avail      0       0       0       0       0       0       0
Rx Avail      0       0       0       0       0       0       0

NP0:
-----
Tx Used       0       0       3000    0       0       0       0
Rx Used       0       0       3000    0       0       0       0
Tx Avail      0       0       0       0       0       0       0
Rx Avail      0       0       0       0       0       0       0

NP1:
-----
Tx Used       0       0       3000    0       0       0       0
Rx Used       0       0       3000    0       0       0       0
Tx Avail      0       0       0       0       0       0       0
Rx Avail      0       0       0       0       0       0       0

NP2:
-----
Tx Used       0       0       0       0       0       0       0
Rx Used       0       0       0       0       0       0       0
Tx Avail      0       0       0       0       0       0       0
Rx Avail      0       0       0       0       0       0       0

NP3:
-----
Tx Used       0       0       2000    0       0       0       0
Rx Used       0       0       2000    0       0       0       0
Tx Avail      0       0       0       0       0       0       0
Rx Avail      0       0       0       0       0       0       0
```

Related Commands

Command	Description
show bfd counters, on page 692	Displays BFD counter information.

show bfd mib session

To display IPv4 and IPv6 Bidirectional Forwarding Detection (BFD) MIB session information, use the **show bfd mib session** command in EXEC mode.

show bfd mib session [**location** *node-id*]

Syntax Description

location *node-id*	(Optional) Displays all IPv4 and IPv6 BFD MIB session information stored on the specified node. The *node-id* argument is entered in the *rack/slot/module* notation.

Command Default

When *node-id* is not specified, information for all IPv4 and IPv6 BFD MIB sessions, stored on the route processor node, is displayed.

Command Modes

EXEC

Command History

Release	Modification
Release 3.9.0	This command was introduced.
Release 4.0.0	Support for the display of IPv6 BFD MIB session information was added.

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.

When *node-id* is not specified, information for all IPv4 and IPv6 BFD MIB sessions, stored on the route processor node, is displayed, and this information is populated and updated only after SNMP operations for those BFD MIB sessions are performed.

When *node-id* is specified, information for all IPv4 and IPv6 BFD MIB sessions, stored on the specified node (linecard), is displayed, and this information is updated automatically without SNMP operations being performed.

Task ID

Task ID	Operations
bgp	read
ospf	read
isis	read
mpls-te	read

Examples

The following example displays all IPv4 and IPv6 BFD MIB session information stored on the RP node:

```
RP/0/RSP0/CPU0:router# show bfd mib session
```

```
Tue Sep  9 07:49:30.828 PST DST
Local Discr: 327681(0x50001), Remote Discr: 0(0x0)
BFD session: GigabitEthernet0_1_5_2(0x11800c0), 10.27.4.7
  Current State: ADMIN DOWN, Number of Times UP: 0
  Running Version: 0, Last Down Diag: None
  Last Up Time (s.ns): 0.0
  Last Down Time (s.ns): 0.0
  Detection Multiplier: 0
  Desired Min TX Interval: 0
  Required Min RX Interval: 0
  Required Min RX Echo Interval: 0
  Packets in/out: 0/0
  Current Trap Bitmap: 0x0
  Last Time Cached: Not yet cached
```

The following example displays all IPv4 and IPv6 BFD MIB session information stored on 0/1/CPU0:

```
RP/0/RSP0/CPU0:router# show bfd mib session location 0/1/CPU0

Tue Sep  9 07:44:49.190 PST DST
Local Discr: 327681(0x50001), Remote Discr: 0(0x0)
BFD session: GigabitEthernet0_1_5_2(0x11800c0), 10.27.4.7
  Number of times UP: 0
  Last Down Diag: None
  Last Up Time (s.ns): 0.0
  Last Down Time (s.ns): 0.0
  Packets in/out: 0/1140134
```

Table 48: show bfd mib Field Descriptions

Field	Description
date and timestamp	Date and time stamp during which a snapshot of the BFD MIB session information is taken.
Local Discr	Local discriminator (in decimal and hexadecimal) that uniquely identifies the BFD MIB session.
Remote Discr	Session discriminator (in decimal and hexadecimal) that was chosen by the remote system for the BFD MIB session.
BFD session	Index of interface upon which the BFD MIB session is running. Also, neighboring IP address that is monitored with the BFD MIB session.
Current State	Current state of the BFD MIB session.
Number of Times UP	Number of times the BFD MIB session has gone into the up state since the router was last rebooted.
Running Version	BFD protocol version number in which the BFD MIB session is running.
Last Down Diag	Diagnostic value associated with the last time the BFD MIB session went down.
Last Up Time (s.ns)	Value of sysUpTime, in *seconds.nanoseconds*, during which the BFD MIB session last came up. If such an event does not exist, a zero is displayed.

Field	Description
Last Down Time (s.ns)	Value of sysUpTime, in *seconds.nanoseconds*, during which communication was last lost with the neighbor. If such an event does not exist, a zero is displayed.
Detection Multiplier	Failure detection multiplier.
Desired Min TX Interval	Minimum interval, in microseconds, preferred by the local system when transmitting BFD control packets.
Required Min RX Interval	Minimum interval, in microseconds, that the local system supports between received BFD control packets.
Required Min RX Echo Interval	Minimum interval, in microseconds, that the local system supports between received BFD echo packets.
Packets in/out	Total number of BFD messages received and transmitted for the BFD MIB session.
Current Trap Bitmap	Bits that control the trap for the BFD MIB session. A nonzero value indicates that the trap is generated when the next trap event is triggered.
Last Time Cached	When information for the BFD MIB session was last cached. Typically, the information is cached when SNMP operations for the BFD MIB session are performed.

Related Commands

Command	Description
show bfd session, on page 702	

show bfd multipath

To display information concerning only BFD multipath sessions, use the **show bfd multipath** command in the EXEC mode.

show bfd multipath {**ipv4** | **ipv6** | **label** | **all**} **location** *node-id*

Syntax Description

ipv4	Displays BFD over IPv4 information only.
ipv6	Displays BFD over IPv6 information only.
label	Displays BFD label information.
all	Displays both BFD over IPv4 and BFD over IPv6 information.
location *node-id*	Displays BFD counters from the specified location. The *node-id* argument is entered in the *rack/slot/module* notation.

Command Default

No default behavior or values

Command Modes

EXEC

Command History

Release	Modification
Release 4.2.0	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
ospf	read
isis	read
mpls-te	read

Example

This example shows the sample output for **show bfd multipath** command:

```
RP/0/RSP0/CPU0:router#show bfd multipath location 0/5/cpu0
```

```
Int/Src Addr    Label/Dest Addr VRF ID     Discr      Node      State
--------------- --------------- ---------- ---------- --------- ---------
pw-ether 1      10.10.10.10     0x00000002 0x4        0/5/CPU0  DOWN
tunnel-ip 1     1.1.1.1         0x8        0x5        0/5/CPU0  UP
```

show bfd session

To display Bidirectional Forwarding Detection (BFD) session information, use the **show bfd session** command in EXEC mode.

show bfd [**ipv4** | [**singlehop** | **multihop**] | **ipv6** | [**singlehop** | **multihop**] | **all** | **label**] **session** [**interface** *type interface-path-id* [**destination** *ip-address*] [**detail**][**in-label**]] **location** *node-id*

Syntax Description

ipv4	(Optional) Displays BFD over IPv4 information only.
ipv6	(Optional) Displays BFD over IPv6 information only.
singlehop	(Optional) Displays BFD singlehop information only.
multihop	(Optional) Displays BFD multihop information only.
all	(Optional) Displays both BFD over IPv4 and BFD over IPv6 information.
label	(Optional) Displays the MPLS Transport Profile (MPLS-TP) label BFD information only.
interface	(Optional) Specifies the interface for which to show information.
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.
destination *ip-address*	(Optional) Displays the BFD session destined for the specified IP address.
detail	(Optional) Displays detailed session information, including statistics and number of state transitions.
in-label	(Optional) Displays the BFD session with a specific incoming MPLS-TP label.
location *node-id*	(Optional) Displays BFD sessions hosted from the specified location. The *node-id* argument is entered in the *rack/slot/module* notation.

Command Default

The default is the default address family identifier (AFI) that is set by the **set default-afi** command, IPv4 or IPv6.

Command Modes

EXEC

Command History

Release	Modification
Release 3.7.2	This command was introduced.

Release	Modification
Release 4.0.0	Support for the **ipv6** keyword was added.
Release 4.2.0	Support for the **singlehop** and **multihop** keywords were added.

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.

For the *interface-path-id* argument, use the following guidelines:

- If specifying a physical interface, the naming notation is *rack*/*slot*/*module*/*port*. The slash between values is required as part of the notation. An explanation of each component of the naming notation is as follows:
  - *rack*: Chassis number of the rack.
  - *slot*: Physical slot number of the line card.
  - *module*: Module number. A physical layer interface module (PLIM) is always 0.
  - *port*: Physical port number of the interface.
- If specifying a virtual interface, the number range varies, depending on interface type.
- The BFD UP time in the show output printed using CLI, and the show output from SNMP may differ. The CLI output is the appropriate value. The difference doesn't impact the functioning of the device.

Note Only VRF ID is displayed in the summary CLI (such as **show bfd multiple-path**, **show bfd all session**, **show bfd counters**) and VRF name and VRF ID is displayed in the detailed CLI (such as **show bfd all session detail**, **show bfd all session status**).

Task ID

Task ID	Operations
bgp	read
ospf	read
isis	read
mpls-te	read

Examples

The following example shows the output from the **show bfd session** command with the **detail** keyword and IPv4 as the default:

```
RP/0/RSP0/CPU0:router# show bfd session detail

I/f:TenGigE0/2/0/0.6, Location:0/2/CPU0, dest:10.0.6.2, src:10.0.6.1
 State:UP for 0d:0h:3m:4s, number of times UP:1
Session type: PR/V4/SH
```

```
Received parameters:
 Version:1, desired tx interval:2 s, required rx interval:2 s
 Required echo rx interval:1 ms, multiplier:3, diag:None
 My discr:589830, your discr:590028, 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:3, diag:None
 My discr:590028, your discr:589830, 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:250 ms, local negotiated echo tx interval:250 ms
 Echo detection time:750 ms(250 ms*3), async detection time:6 s(2 s*3)
Local Stats:
 Intervals between async packets:
   Tx:Number of intervals=100, min=952 ms, max=2001 ms, avg=1835 ms
       Last packet transmitted 606 ms ago
   Rx:Number of intervals=100, min=1665 ms, max=2001 ms, avg=1828 ms
       Last packet received 1302 ms ago
 Intervals between echo packets:
   Tx:Number of intervals=100, min=250 ms, max=252 ms, avg=250 ms
       Last packet transmitted 188 ms ago
   Rx:Number of intervals=100, min=250 ms, max=252 ms, avg=250 ms
       Last packet received 187 ms ago
 Latency of echo packets (time between tx and rx):
   Number of packets:100, min=1 ms, max=2 ms, avg=1 ms
                           Desired               Adjusted
  Client               Interval   Multiplier Interval   Multiplier
  -------------------- ---------- ---------- ---------- ----------
  ipv4_static          500 ms     3          500 ms     3
  bgp-default          1500 ms    3          1500 ms    3
```

The following example shows the output from the **show bfd session** command with the **all** keyword, which displays both IPv4 and IPv6 information:

```
RP/0/RSP0/CPU0:router# show bfd all session location 0/1/CPU0

Mon Nov  5 08:51:50.339 UTC
IPv4:
-----
Interface            Dest Addr           Local det time(int*mult)      State
                                           Echo            Async
-------------------- --------------- ---------------- ---------------- ---------
PO0/1/0/0            10.0.0.2        300ms(100ms*3)   6s(2s*3)         UP

IPv6:
-----
Interface            Dest Addr
                     Local det time(int*mult)          State
                            Echo            Async
------------------- ------------------------------------------------
PO0/1/0/0            abcd::2
                     0s(0s*0)         15s(5s*3)        UP
```

Table 49: show bfd session detail command Field Descriptions

Field	Description
I/f	Interface type.

Field	Description
Location	Location of the node that hosts the local endpoint of the connection, in the *rack/slot/module* notation
dest	IP address of the destination endpoint.
src	IP address of the source endpoint.
State	Current state of the connection, and the number of days, hours, minutes, and seconds that this connection has been active.
number of times UP	Number of times this connection has been brought up.
Received parameters	Provides information on the last transmitted control packet for the session: • Version—Version number of the BFD protocol. • desired tx interval—Desired transmit interval. • required rx interval—Required receive interval. • Required echo rx interval—Required echo receive interval. • multiplier— Number of times a packets is missed before BFD declares the neighbor down. • diag—diagnostic code specifying the peer system's reason for the last transition of the session from Up to some other state. • My discr—unique, nonzero discriminator value generated by the transmitting system, used to demultiplex multiple BFD sessions between the same pair of systems. • your discr— discriminator received from the corresponding remote system. This field reflects back the received value of My discr, or is zero if that value is unknown.
Transmitted parameters	Provides information on the last transmitted control packet for the session: • Version—Version number of the BFD protocol. • desired tx interval—Desired transmit interval. • required rx interval—Required receive interval • Required echo rx interval—Required echo receive interval • multiplierNumber of times a packets is missed before BFD declares the neighbor down. • diag—diagnostic code specifying the local system's reason for the last transition of the session from Up to some other state. • My discr—unique, nonzero discriminator value generated by the transmitting system, used to demultiplex multiple BFD sessions between the same pair of systems. • your discr— discriminator received from the corresponding remote system. This field reflects back the received value of My discr, or is zero if that value is unknown.

Field	Description
Timer Values	Provides information on the timer values used by the local and remote ends, as follows: • Local negotiated async tx interval—interval at which control packets are being transmitted by the local end. • Remote negotiated async tx interval—interval at which control packets should be transmitted by the remote end. • Desired echo tx interval—interval at which the local end would like to transmit echo packets. • local negotiated echo tx interval—interval at which echo packets are being transmitted by the local end. • Echo detection time—local failure detection time of echo packets. It is the product of the local negotiated echo tx interval and the local multiplier. • async detection time—local failure detection time of the asynchronous mode (control packets). It is the product of the remote negotiated async tx interval and the remote multiplier.

Local Stats	Displays the local transmit and receive statistics,
	• Intervals between async packets—provides measurements on intervals between control packets (tx and rx):
	  • Number of intervals—number of sampled intervals between control packets
	  • min—minimum measured interval between 2 consecutive control packets
	  • max—maximum measured interval between 2 consecutive control packets
	  • avg—average measured interval between 2 consecutive control packets
	  • Last packet received/transmitted—indicates how long ago the last control packet was received/transmitted.
	• Intervals between echo packets—provides measurements on intervals between echo packets (tx and rx). The measurements have the same meaning as for async packets.
	• Latency of echo packets (time between tx and rx)—provides measurements on latency of echo packets, i.e. the time between tx and rx of echo packets:
	  • Number of packets—number of sampled echo packets.
	  • min—minimum measured latency for echo packets.
	  • max—maximum measured latency of echo packets.
	  • avg—average measured latency of echo packets.

Field	Description
Session owner information	Provides the following information about the session owner. • Client—name of the client application process. • Desired interval—desired interval provided by the client, in milliseconds. • Multiplier—multiplier value provided by the client.

Related Commands

Command	Description
bfd fast-detect, on page 652	Enables BFD to detect failures in the path between adjacent forwarding engines.
bfd multiplier, on page 662	Sets the BFD multiplier.
show bfd mib session, on page 697	Displays BFD MIB session information.

show bfd summary

To display the percentage of PPS rate in use per line card, maximum usage of PPS, and total number of sessions, use the **show bfd summary** command in the EXEC mode.

show bfd summary [**private**]**location***node-id*

Syntax Description

private	Displays the private information.
location *node-id*	Displays BFD counters from the specified location. The *node-id* argument is entered in the *rack/slot/module* notation.

Command Default

No default behavior or values

Command Modes

EXEC

Command History

Release	Modification
Release 4.2.0	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
ospf	read
isis	read
mpls-te	read

Example

This example shows the sample output from the **show bfd summary** command for a specified location:

```
RP/0/RSP0/CPU0:router#show bfd summary location 0/1/cpu0

Node       PPS rate usage  Session number
           %   Used  Max   Total   Max
---------- --------------- --------------
0/1/CPU0   0   80    9600  4       4000
```

This example shows the sample output from the **show bfd summary** command:

```
RP/0/RSP0/CPU0:router#show bfd summary
Node       PPS rate usage  Session number
           %   Used  Max   Total   Max
---------- --------------- --------------
0/0/CPU0   0   0     9600  0       4000
0/1/CPU0   0   0     9600  0       4000
0/2/CPU0   0   0     9600  0       4000
0/5/CPU0   0   0     9600  0       4000
0/6/CPU0   0   0     9600  0       4000
0/7/CPU0   0   0     9600  0       4000
```



EIGRP Commands

This module describes the commands used to configure and monitor the Enhanced Interior Gateway Routing Protocol (EIGRP).

For detailed information about EIGRP concepts, configuration tasks, and examples, see *Implementing EIGRP on* Cisco ASR 9000 Series Router in *Routing Configuration Guide for Cisco ASR 9000 Series Routers*.

- address-family (EIGRP), on page 713
- authentication keychain, on page 715
- auto-summary (EIGRP), on page 717
- autonomous-system, on page 719
- bandwidth-percent (EIGRP), on page 720
- bfd fast-detect (EIGRP), on page 721
- bfd minimum-interval (EIGRP), on page 722
- bfd multiplier (EIGRP), on page 724
- clear eigrp neighbors, on page 725
- clear eigrp topology, on page 727
- default-information, on page 729
- default-metric (EIGRP), on page 730
- distance (EIGRP), on page 732
- hello-interval (EIGRP), on page 734
- hold-time (EIGRP), on page 735
- interface (EIGRP), on page 737
- log-neighbor-changes, on page 739
- log-neighbor-warnings, on page 740
- maximum-paths (EIGRP), on page 741
- maximum-prefix (EIGRP), on page 742
- metric (EIGRP), on page 744
- metric maximum-hops, on page 746
- metric rib-scale, on page 747
- metric version (EIGRP), on page 748
- metric weights, on page 749
- neighbor (EIGRP), on page 751
- neighbor maximum-prefix, on page 753
- next-hop-self disable, on page 756
- nsf disable (EIGRP), on page 757

- passive-interface (EIGRP), on page 759
- redistribute (EIGRP), on page 760
- redistribute maximum-prefix, on page 762
- remote-neighbor (unicast), on page 764
- route-policy (EIGRP), on page 766
- router eigrp, on page 767
- router-id (EIGRP), on page 768
- show eigrp accounting, on page 769
- show eigrp interfaces, on page 771
- show eigrp neighbors, on page 775
- show eigrp topology, on page 778
- show eigrp traffic, on page 781
- show protocols (EIGRP), on page 783
- site-of-origin (EIGRP), on page 786
- split-horizon disable (EIGRP), on page 788
- stub (EIGRP), on page 789
- summary-address (EIGRP), on page 791
- timers active-time, on page 793
- timers nsf route-hold (EIGRP), on page 794
- variance, on page 795
- vrf (EIGRP), on page 796

address-family (EIGRP)

To enable an IPv4 address family under Enhanced Interior Gateway Routing Protocol (EIGRP), use the **address-family** command in the appropriate mode. To remove the address family from the EIGRP configuration, use the **no** form of this command.

address-family {**ipv4** | **ipv6**}
no address-family {**ipv4** | **ipv6**}

Syntax Description

ipv4	Selects IPv4 address family.

Command Default

No default behavior or values

Command Modes

Router configuration

VRF 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **address-family ipv4** command to configure IPv4 address family sessions in EIGRP.

EIGRP Virtual Private Networks (VPNs) can be configured under IPv4 address family configuration modes after first entering VRF configuration mode. All commands in address family configuration mode can be configured in VRF address families except the **autonomous-system** and **maximum-prefix** commands.

Task ID

Task ID	Operations
eigrp	read, write

Examples

The following example shows how to configure an IPv4 VRF address family session after defining the VRF named vrf1:

```
RP/0/RSP0/CPU0:router(config)# router eigrp 1
RP/0/RSP0/CPU0:router(config-eigrp)# vrf vrf1
RP/0/RSP0/CPU0:router(config-eigrp-vrf)# address-family ipv4
RP/0/RSP0/CPU0:router(config-eigrp-vrf-af)# default-metric 1000 100 255 1 1500
```

Related Commands

Command	Description
autonomous-system	Configures an EIGRP routing process to run within a VRF.
maximum-prefix (EIGRP)	Limits the number of prefixes that are accepted under an address family by EIGRP.

authentication keychain

To authenticate all EIGRP protocol traffic on one or more interfaces based on the MD5 algorithm, use the **authentication keychain** command in an appropriate configuration mode. To disable authentication, use the **no** form of this command.

authentication keychain *key-chain-name*
no authentication keychain *key-chain-name*

Syntax Description

key-chain-name	Name of the authentication keychain

Command Default

Authentication is disabled.

Command Modes

IPv4 address family interface configuration

IPv6 address family interface configuration

IPv4 VRF address family interface configuration

IPv6 VRF address family 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **address-family ipv4** command to configure IPv4 address family sessions in EIGRP and the **address-family ipv6** command to configure IPv6 address family sessions in EIGRP.

EIGRP Virtual Private Networks (VPNs) can be configured under IPv4 and IPv6 address family configuration modes after first entering VRF configuration mode.

Task ID

Task ID	Operations
eigrp	read, write

Examples

The following example shows how to enable an EIGRP authentication keychain:

```
RP/0/RSP0/CPU0:router# configure eigrp 1
RP/0/RSP0/CPU0:router(config-eigrp)# address-family ipv4
RP/0/RSP0/CPU0:router(config-eigrp-af)# interface GigabitEthernet 0/2/0/0
RP/0/RSP0/CPU0:router(config-eigrp-if)# authentication key chain key1
```

Related Commands

Command	Description
router eigrp, on page 767	Configures a routing process and enter router configuration mode for Enhanced Interior Gateway Routing Protocol (EIGRP).

auto-summary (EIGRP)

To allow automatic summarization of subnet routes into network-level routes for an Enhanced Interior Gateway Routing Protocol (EIGRP) process, use the **auto-summary** command in the appropriate configuration mode. To disable this function and send subprefix routing information across classful network boundaries, use the **no** form of this command.

auto-summary
no auto-summary

Syntax Description

This command has no keywords or arguments.

Command Default

The behavior of this command is disabled by default. (The software sends subnet routing information across classful network boundaries.)

Command Modes

IPv4 Address family configuration

IPv4 VRF address family 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Route summarization reduces the amount of routing information in the routing tables. Use the **auto-summary** command to allow the software to create summary subprefixes to the classful network boundary when crossing classful network boundaries.

EIGRP summary routes are given an administrative distance value of 5. You cannot configure this value.

Task ID

Task ID	Operations
eigrp	read, write

Examples

The following example shows how to enable automatic summarization for EIGRP 1:

```
RP/0/RSP0/CPU0:router(config)# router eigrp 1
RP/0/RSP0/CPU0:router(config-eigrp)# vrf vpn-1

RP/0/RSP0/CPU0:router(config-eigrp-vrf)# address-family ipv4
RP/0/RSP0/CPU0:router(config-eigrp-vrf-af)# auto-summary
```

Related Commands

Command	Description
summary-address (EIGRP), on page 791	Configures a summary aggregate address for a specified interface.

autonomous-system

To configure an Enhanced Interior Gateway Routing Protocol (EIGRP) routing process to run within a VPN routing and forwarding instance (VRF), use the **autonomous-system** command in IPv4 VRF address family configuration mode.

autonomous-system *as-number*
no autonomous-system *as-number*

Syntax Description

as-number	Autonomous system number of the EIGRP routing process. Range is from 1 to 65535.

Command Default

None

Command Modes

IPv4 VRF address family 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **autonomous-system** command in IPv4 VRF address family configuration mode.

Task ID

Task ID	Operations
eigrp	read, write

Examples

This example shows how to configure autonomous system 101 under VRF VPN-1:

```
RP/0/RSP0/CPU0:router(config)# router eigrp 1
 RP/0/RSP0/CPU0:router(config-eigrp)# vrf vpn-1
RP/0/RSP0/CPU0:router(config-eigrp-vrf)# address-family ipv4
RP/0/RSP0/CPU0:router(config-eigrp-vrf-af)# autonomous-system 101
```

Related Commands

Command	Description
vrf (EIGRP), on page 796	Defines a VRF instance and to enters VRF configuration mode.

bandwidth-percent (EIGRP)

To configure the percentage of bandwidth that may be used by EIGRP on an interface, use the **bandwidth-percent** command in interface configuration mode. To restore the default value, use the **no** form of this command.

bandwidth-percent *percent*
no bandwidth-percent

Syntax Description

percent	Percentage of bandwidth that EIGRP may use.

Command Default

percent : 50

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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

EIGRP uses up to 50 percent of the bandwidth of a link, as defined by the bandwidth interface configuration command. This command may be used if some other fraction of the bandwidth is desired. Values greater than 100 percent may be configured. The configuration option may be useful if the bandwidth is set artificially low for other reasons.

Task ID

Task ID	Operations
eigrp	read, write

Examples

The following example shows how to configure EIGRP to use up to 75 percent (42 kbps) of an interface in autonomous system 209:

```
RP/0/RSP0/CPU0:router(config)# router eigrp 1
RP/0/RSP0/CPU0:router(config-eigrp)# address-family ipv4
RP/0/RSP0/CPU0:router(config-eigrp-af)# router-id 10.1.1.1
RP/0/RSP0/CPU0:router(config-eigrp-af)# interface GigabitEthernet 0/1/0/0
RP/0/RSP0/CPU0:router(config-eigrp-af-if)# bandwidth-percent 75
```

Related Commands

Command	Description
bandwidth (interface)	Sets a bandwidth value for an interface.

bfd fast-detect (EIGRP)

To enable Bidirectional Forwarding Detection (BFD) to detect failures in the path between adjacent forwarding engines, use the **bfd fast-detect** command in router configuration mode. To return the software to the default state in which BFD is not enabled, use the **no** form of this command.

bfd fast-detect
no bfd fast-detect

Syntax Description

This command has no keywords or arguments.

Command Default

BFD detection of failures in the path between adjacent forwarding engines is disabled.

Command Modes

Router configuration

Command History

Release	Modification
Release 5.2.0	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
eigrp	read, write

Examples

The following example shows how to configure BFD on a EIGRP router:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router eigrp 100
RP/0/RSP0/CPU0:router(config-eigrp)# address-family ipv4
RP/0/RSP0/CPU0:router(config-eigrp-af)# interface gigabitEthernet 2/2/2/1
RP/0/RSP0/CPU0:router(config-eigrp-af-if)# bfd fast-detect
```

bfd minimum-interval (EIGRP)

To specify the minimum control packet interval for BFD sessions for the corresponding BFD configuration scope, use the **bfd minimum-interval** command in router configuration mode. To return the router to the default setting, use the **no** form of this command.

bfd minimum-interval *milliseconds*
no bfd minimum-interval [*milliseconds*]

Syntax Description

milliseconds	Interval between sending BFD hello packets to the neighbor. The range is 15 to 30000 milliseconds.

Command Default

EIGRP *interval*: 50 milliseconds

Command Modes

Router configuration

Command History

Release	Modification
Release 5.2.0	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
eigrp	read, write

Examples

The following example shows how to configure BFD minimum interval on a EIGRP router:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router eigrp 100
RP/0/RSP0/CPU0:router(config-eigrp)# address-family ipv4
RP/0/RSP0/CPU0:router(config-eigrp-af)# interface gigabitEthernet 2/2/2/1
RP/0/RSP0/CPU0:router(config-eigrp-af-if)# bfd minimum-interval 50
```

Related Commands

Command	Description
address-family ipv4 unicast (BFD), on page 627	Enables BFD fast-detection on a specific IPV4 unicast destination address prefix and on the forwarding next-hop address.
bfd minimum-interval, on page 655	Specifies the minimum control packet interval for BFD sessions for the corresponding BFD configuration scope.
bfd multiplier, on page 662	Sets the BFD multiplier.

Command	Description
show bfd, on page 687	Displays BFD information for a specific location.

bfd multiplier (EIGRP)

To set the Bidirectional Forwarding Detection (BFD) multiplier, use the **bfd multiplier** command in the appropriate configuration mode. To return the router to the default setting, use the **no** form of this command.

bfd multiplier *multiplier*
no bfd multiplier [*multiplier*]

Syntax Description

multiplier	Number of times a packet is missed before BFD declares the neighbor down. The range is 2 to 50.

Command Default

The default multiplier is 3.

Command Modes

Router configuration

Command History

Release	Modification
Release 5.2.0	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
eigrp	read, write

Examples

The following example shows how to configure BFD minimum interval on a EIGRP router:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router eigrp 100
RP/0/RSP0/CPU0:router(config-eigrp)# address-family ipv4
RP/0/RSP0/CPU0:router(config-eigrp-af)# interface gigabitEthernet 2/2/2/1
RP/0/RSP0/CPU0:router(config-eigrp-af-if)# bfd multiplier 5
```

clear eigrp neighbors

To remove and re-establish Enhanced Interior Gateway Routing Protocol (EIGRP) neighbor entries from the appropriate table, use the **clear eigrp neighbors** command in EXEC configuration mode.

clear eigrp [*as-number*] [**vrf** {**vrf***all*}] [**ipv4** | **ipv6**] **neighbors** [*ip-addresstype interface-path-id*] [**soft**]

Syntax Description

as-number	(Optional) Autonomous system number. Range is from 1 to 65535.
vrf { *vrf* | **all** }	(Optional) Specifies a particular VPN routing and forwarding instance (VRF) or all VRF instances.
ipv4	(Optional) Specifies the IPv4 address family.
ip-address	(Optional) Address of the neighbor.
type	Interface type. For more information, use the question mark (**?**) online help function.
interface-path-id	Physical interface or virtual interface. 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.
soft	(Optional) Specifies a soft reset.

Command Default

When no autonomous system number or VRF instance is specified, all EIGRP neighbor entries are cleared from the table.

Command Modes

EXEC 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 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
eigrp	read, write

Examples

The following example shows how to clear all EIGRP VRF entries for neighbor Gigabit Ethernet interface 0/5/0/0:

```
RP/0/RSP0/CPU0:router# clear eigrp customer_1 neighbors GigabitEthernet 0/5/0/0
```

Related Commands

Command	Description
show eigrp interfaces, on page 771	Displays information about interfaces configured for EIGRP.
show eigrp neighbors, on page 775	Displays the neighbors discovered by EIGRP.

clear eigrp topology

To remove and relearn Enhanced Interior Gateway Routing Protocol (EIGRP) topology entries from the appropriate table, use the **clear eigrp topology** command in EXEC configuration mode.

clear eigrp [*as-number*] [**vrf** {**vrf***all*}] [**ipv4** | **ipv6**] **topology** [*prefix mask prefixl/length*]

Syntax Description

as-number	(Optional) Autonomous system number. Range is from 1 to 65535.
vrf { *vrf* | **all** }	(Optional) Specifies a particular VPN routing and forwarding instance (VRF) or all VRF instances.
ipv4	(Optional) Specifies the IPv4 address family.
prefix	IP prefix, which limits output to a specific route.
mask	IP address mask.
/ *length*	Prefix length, which can be indicated as a slash (/) and number. For example, /8 indicates that the first eight bits in the IP prefix are network bits. If *length* is used, the slash is required.

Command Default

No EIGRP topology entries are cleared.

Command Modes

EXEC 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 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
eigrp	read, write

Examples

The following example shows how to clear EIGRP topology entries for a specific route:

```
RP/0/RSP0/CPU0:router# clear eigrp topology 10.1.0.0/8
```

Related Commands

Command	Description
show eigrp topology, on page 778	Displays information for the EIGRP topology table.

default-information

To control the candidate default routing information for an Enhanced Interior Gateway Routing Protocol (EIGRP), use the **defaultinformation** command in the appropriate configuration mode. To suppress EIGRP candidate default information in incoming or outgoing updates, use the **no** form of this command.

default-information allowed {**in** | **out**} [**route-policy** *name*]
no default-information allowed {**in** | **out**} [**route-policy** *name*]

Syntax Description

allowed	Specifies EIGRP to allow default routing information.
in	Specifies EIGRP to allow inbound default routing information.
out	Specifies EIGRP to allow outbound default routing information.
route-policy *name*	(Optional) Specifies a route policy.

Command Default

Default routing information is not accepted or flagged.

Command Modes

Address family configuration

IPv4 VRF address family 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 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
eigrp	read, write

Examples

The following example shows how to configure inbound default routes specified with route policy acme to be accepted by an EIGRP peer in autonomous system 1:

```
RP/0/RSP0/CPU0:router(config)# router eigrp 1
RP/0/RSP0/CPU0:router(config-eigrp)# vrf vrf1
RP/0/RSP0/CPU0:router(config-eigrp-vrf)# address-family ipv4
RP/0/RSP0/CPU0:router(config-eigrp-vrf-af)# default-information accept in route-policy
acme
```

default-metric (EIGRP)

To set metrics for an Enhanced Interior Gateway Routing Protocol (EIGRP), use the **defaultmetric** command in the appropriate configuration mode. To remove the metric values and restore the default state, use the **no** form of this command.

default-metric *bandwidth delay reliability loading mtu*
no default-metric

Syntax Description

bandwidth	Minimum bandwidth of the route in kilobits per second. Range is 1 to 4294967295.
delay	Route delay in ten microsecond units. Range is 1 to 4294967295.
reliability	Likelihood of successful packet transmission expressed as a number between 0 and 255. The value 255 means 100-percent reliability; 0 means the link is not reliable.
loading	Effective bandwidth of the route expressed as a number from 1 to 255 (255 is 100-percent loading).
mtu	Minimum maximum transmission unit (MTU) size of the route in bytes. Range is from 1 to 65535.

Command Default

No default values

Command Modes

IPv4 address family configuration

IPv4 VRF address family 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **default-metric** command to provide default metric values while redistributing a protocol into EIGRP.

Metric defaults have been carefully set to work for a wide variety of networks. Take great care when changing these values.

Task ID

Task ID	Operations
eigrp	read, write

Examples

The following example shows how to take redistributed Routing Information Protocol (RIP) metrics and translate them into EIGRP metrics with values as follows: bandwidth = 1000, delay = 100, reliability = 250, loading = 100, and MTU = 1500.

```
RP/0/RSP0/CPU0:router(config)# router eigrp 1
RP/0/RSP0/CPU0:router(config-eigrp)# vrf vrf1
RP/0/RSP0/CPU0:router(config-eigrp-vrf)# address-family ipv4
RP/0/RSP0/CPU0:router(config-eigrp-vrf-af)# redistribute rip
RP/0/RSP0/CPU0:router(config-eigrp-vrf-af)# default-metric 1000 100 250 100 1500
```

Related Commands

Command	Description
redistribute	Redistributes routes from one routing domain into another routing domain.

distance (EIGRP)

To allow the use of one of two administrative distances—internal and external—that could provide a better route to a node, use the **distance** command in the appropriate configuration mode. To reset these values to their defaults, use the **no** form of this command.

distance *internal-distance external-distance*
no distance

Syntax Description

internal-distance	Administrative distance for EIGRP internal routes. Internal routes are those that are learned from another entity within the same autonomous system (AS). The distance can be a value from 1 to 255.
external-distance	Administrative distance for EIGRP external routes. External routes are those for which the best path is learned from a source external to the AS. The distance can be a value from 1 to 255.

Command Default

internal-distance : 90

external-distance : 170

Command Modes

IPv4 address family configuration

IPv4 VRF address family 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

An administrative distance is a rating of the trustworthiness of a routing information source, such as an individual router or a group of routers. Numerically, an administrative distance is an integer from 0 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.

Use the **distance** command if another protocol is known to provide a better route to a node than was actually learned through the external Enhanced Interior Gateway Routing Protocol (EIGRP) or some internal routes should be preferred by EIGRP.

To display the default administrative distance for a specified routing process, use the **show protocols** EXEC command.

Task ID

Task ID	Operations
eigrp	read, write

Examples

The following example shows how to set the administrative distance of all EIGRP 1 internal routes (within vrf vpn-1) to 80 and all EIGRP external routes to 130:

```
RP/0/RSP0/CPU0:router(config)# router eigrp 1
RP/0/RSP0/CPU0:router(config-eigrp)# vrf vrf1
RP/0/RSP0/CPU0:router(config-eigrp-vrf)# address-family ipv4
RP/0/RSP0/CPU0:router(config-eigrp-vrf-af)# distance 80 130
```

Related Commands

Command	Description
show protocols (EIGRP), on page 783	Displays information about the Enhanced Interior Gateway Routing Protocol (EIGRP) running on the router.

hello-interval (EIGRP)

To configure the hello interval for an interface, use the **hello-interval** command in interface configuration mode. To restore the default value, use the **no** form of this command.

hello-interval *seconds*
no hello-interval

Syntax Description

seconds	Hello interval (in seconds). Range is from 1 to 65535.

Command Default

For low-speed, nonbroadcast multiaccess (NBMA) networks: 60 seconds

For all other networks: 5 seconds

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 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
eigrp	read, write

Examples

The following example shows how to set the hello interval to 10 seconds for the interface:

```
RP/0/RSP0/CPU0:router(config)# router eigrp 1
RP/0/RSP0/CPU0:router(config-eigrp)# address-family ipv4
RP/0/RSP0/CPU0:router(config-eigrp)# router-id 10.1.1.1
RP/0/RSP0/CPU0:router(config-eigrp-af)# interface GigabitEthernet 0/1/0/0
RP/0/RSP0/CPU0:router(config-eigrp-af-if)# hello-interval 10
```

hold-time (EIGRP)

To configure the hold time for an interface, use the **hold-time** command in interface configuration mode. To restore the default value, use the **no** form of this command.

hold-time *seconds*
no hold-time

Syntax Description

seconds	Hold time (in seconds). Range is from 1 to 65535.

Command Default

Three times the default hello interval time of 15 seconds.

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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

On very congested and large networks, the default hold time might not be sufficient time for all routers to receive hello packets from their neighbors. In this case, you may want to increase the hold time.

We recommend that the hold time be at least three times the hello interval. If a router does not receive a hello packet within the specified hold time, routes through this router are considered unavailable.

Increasing the hold time delays route convergence across the network.

To ensure nonstop forwarding during RP failovers, we recommend that the hold time be increased to 30 seconds.

Task ID

Task ID	Operations
eigrp	read, write

Examples

The following example shows how to set the hold time to 0 to 40 seconds for the interface:

```
RP/0/RSP0/CPU0:router(config)# router eigrp 1
RP/0/RSP0/CPU0:router(config-eigrp)# address-family ipv4
RP/0/RSP0/CPU0:router(config-eigrp)# router-id 10.1.1.1
RP/0/RSP0/CPU0:router(config-eigrp-af)# interface GigabitEthernet 0/1/0/0
RP/0/RSP0/CPU0:router(config-eigrp-af-if)# hold-time 40
```

Related Commands

Command	Description
bandwidth (interface)	Sets a bandwidth value for an interface.
hello-interval (EIGRP), on page 734	Configures the hello interval for the EIGRP routing process designated by an autonomous system number.

interface (EIGRP)

To define the interfaces on which the Enhanced Interior Gateway Routing Protocol (EIGRP) routing protocol runs, use the **interface** command in the appropriate configuration mode. To disable EIGRP routing for interfaces, use the **no** form of this command.

interface *type interface-path-id*
no interface *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

When you do not specify this command in configuration mode, EIGRP routing for interfaces is not enabled.

Command Modes

IPv4 address family configuration

IPv4 VRF address family 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **interface** command to associate a specific interface with an EIGRP process. The interface remains associated with the process even when the IPv4 address of the interface changes.

This command places the router in interface configuration mode, from which you can configure interface-specific settings. Commands configured under this mode (such as the **hello-interval** command) are automatically bound to that interface.

Task ID

Task ID	Operations
eigrp	read, write

Examples

The following example shows how to enter interface configuration mode for EIGRP process 1 and set the hello interval to 10 seconds for GigabitEthernet interface 0/1/0/0:

```
RP/0/RSP0/CPU0:router(config)# router eigrp 1
RP/0/RSP0/CPU0:router(config-eigrp)# address-family ipv4
RP/0/RSP0/CPU0:router(config-eigrp)# router-id 10.1.1.1
RP/0/RSP0/CPU0:router(config-eigrp-af)# interface GigabitEthernet 0/1/0/0
RP/0/RSP0/CPU0:router(config-eigrp-af-if)# hello-interval 10
```

log-neighbor-changes

To enable the logging of changes in Enhanced Interior Gateway Routing Protocol (EIGRP) neighbor adjacencies, use the **log-neighbor-changes** command in the appropriate configuration mode. To disable the logging of changes in EIGRP neighbor adjacencies, use the **no** form of this command.

log-neighbor-changes
no log-neighbor-changes

Command Default

This command has no keywords or arguments.

Adjacency changes are not logged.

Command Modes

IPv4 address family configuration

IPv4 VRF address family 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **log-neighbor-changes** command to log neighbor adjacency changes, monitor the stability of the routing system, and help detect problems. Logging is disabled by default. To disable the logging of neighbor adjacency changes, use the **no** form of this command.

Task ID

Task ID	Operations
eigrp	read, write

Examples

The following example shows how to enable logging of neighbor changes for EIGRP 1:

```
RP/0/RSP0/CPU0:router(config)# router eigrp 1
RP/0/RSP0/CPU0:router(config-eigrp)# address-family ipv4
RP/0/RSP0/CPU0:router(config-eigrp-af)# log-neighbor-changes
```

log-neighbor-warnings

To enable the logging of Enhanced Interior Gateway Routing Protocol (EIGRP) neighbor warning messages, use the **log-neighbor-warnings** command in the appropriate configuration mode. To disable the logging of EIGRP neighbor warning messages, use the **no** form of this command.

log-neighbor-warnings
no log-neighbor-warnings

Syntax Description

This command has no keywords or arguments.

Command Default

Neighbor warning messages are not logged.

Command Modes

IPv4 address family configuration

IPv4 VRF address family 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **log-neighbor-warnings** command to disable and enable neighbor warning messages. When neighbor warning messages occur, they are not logged by default.

Task ID

Task ID	Operations
eigrp	read, write

Examples

The following example shows how to configure log neighbor warning messages for EIGRP process 20:

```
RP/0/RSP0/CPU0:router(config)# router eigrp 20
RP/0/RSP0/CPU0:router(config-eigrp) vrf vrf1
RP/0/RSP0/CPU0:router(config-eigrp-vrf)# address-family ipv4
RP/0/RSP0/CPU0:router(config-eigrp-vrf-af)# log-neighbor-warnings
```

maximum-paths (EIGRP)

To control the maximum number of parallel routes that the Enhanced Interior Gateway Routing Protocol (EIGRP) can support, use the **maximum-paths** command in the appropriate configuration mode. To remove the **maximum-paths** command from the configuration file and restore the system to its default condition with respect to the routing protocol, use the **no** form of this command.

maximum-paths *maximum*
no maximum-paths

Syntax Description

maximum	Maximum number of parallel routes that EIGRP can install in a routing table. Range is from 1 to 32 routes .

Command Default

maximum : 4

Command Modes

IPv4 address family configuration

IPv4 VRF address family 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **maximum-paths** command to allow the EIGRP protocol to install multiple paths into the routing table for each prefix. Multiple paths are installed for both internal and external routes, providing these routes are learned in the same autonomous system and that they are equal cost (according to the EIGRP best path algorithm).

Task ID

Task ID	Operations
eigrp	read, write

Examples

The following example shows how to allow a maximum of 10 paths to a destination:

```
RP/0/RSP0/CPU0:router(config)# router eigrp 1
RP/0/RSP0/CPU0:router(config-eigrp) vrf vrf1
RP/0/RSP0/CPU0:router(config-eigrp-vrf)# address-family ipv4
RP/0/RSP0/CPU0:router(config-eigrp-vrf-af)# maximum-paths 10
```

maximum-prefix (EIGRP)

To limit the number of prefixes that are accepted under a VRF address family by Enhanced Interior Gateway Routing Protocol (EIGRP), use the **maximum-prefix** command in IPv4 VRF address family configuration mode. To disable this function, use the **no** form of this command.

maximum-prefix *maximum* [*threshold*] [**dampened**] [**reset-time** *minutes*] [**restart** *minutes*] [**restart-count** *number*] [**warning-only**]
no maximum-prefix *maximum* [*threshold*] [**dampened**] [**reset-time** *minutes*] [**restart** *minutes*] [**restart-count** *number*] [**warning-only**]

Syntax Description

maximum	Maximum number of prefixes allowed under an address family. Range is from 1 to 4294967295. The number of prefixes that can be configured is limited only by the available system resources on the router.
threshold	(Optional) Syslog warning messages are specified as a percentage of the maximum prefix limit that was exceeded. The prefix percentage number range is from 1 to 100. The default is 75 percent.
dampened	(Optional) A decay penalty is applied to the restart time period each time the maximum prefix limit is exceeded. The half-life for the decay penalty is 150 percent of the default or user-defined restart time value in minutes. This keyword is disabled by default.
reset-time *minutes*	(Optional) The restart count is reset to 0 after the default or user-defined reset time period has expired. The range of values that can be applied with the *minutes* argument is from 1 to 65535 minutes. The default reset-time period is 15 minutes.
restart *minutes*	(Optional) A time period when router adjacencies are not formed or when redistributed routes are not accepted from the RIB after the maximum prefix limit has been exceeded. The value for the *minutes* argument is from 1 to 65535 minutes. The default restart time period is 5 minutes.
restart-count *number*	(Optional) Number of times a peering session is automatically reestablished after the peering session is torn down or after the redistribute route is cleared and relearned when the maximum prefix exceeds limits. The default restart count limit is 3. **Caution** After the restart count threshold is crossed, you need to use the `clear eigrp neighbors` command to re-establish normal peering, redistribution, or both.
warning-only	(Optional) Configures the router to generate syslog messages only when the maximum prefix limit is reached, instead of terminating the peering session.

Command Default

threshold: 75 percent

dampened : False

reset-time : 15 minutes

restart : 5 minutes

restart-count : 3

warning-only : False

Command Modes

IPv4 VRF address family 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **maximum-prefix** command to limit the number of prefixes that are accepted from all sources. When the maximum prefix limit is exceeded, sessions with remote peers are torn down, all routes learned from remote peers and through redistribution are removed from the topology and routing tables, and redistribution and peering are suspended for the default or user-defined time period.

Task ID

Task ID	Operations
eigrp	read, write

Examples

The following example shows how to configure the maximum prefix limit for an EIGRP process, which includes routes learned through redistribution and routes learned through EIGRP peering sessions. The maximum limit is set to 50,000 prefixes. When the number of prefixes learned through redistribution reaches 37,500 (75 percent of 50,000), warning messages are displayed in the console. When the maximum prefix limit is exceeded, all peering sessions are reset, the topology and routing tables are cleared and redistributed routes and all peering sessions are placed in a penalty state.

```
RP/0/RSP0/CPU0:router(config)# router eigrp 1
RP/0/RSP0/CPU0:router(config-eigrp)# vrf vrf1
RP/0/RSP0/CPU0:router(config-eigrp-vrf)# address-family ipv4
RP/0/RSP0/CPU0:router(config-eigrp-vrf-af)# maximum-prefix 50000
```

Related Commands

Command	Description
clear eigrp neighbors, on page 725	Deletes EIGRP VPN neighbor entries from the table.
clear route	Deletes routes from the IP routing table.

metric (EIGRP)

To set metrics for an Enhanced Interior Gateway Routing Protocol (EIGRP) interface, use the **metric** command in interface configuration mode. To remove the metric values and restore the default state, use the **no** form of this command.

metric {**bandwidth** | **delay** | **load** | **reliability**}
no metric

Syntax Description

bandwidth	Minimum interface bandwidth of the route in kilobits per second. Range is 1 to 4294967295.
delay	Interface route delay in tens of microseconds. Delay is 1 or any positive number that is a multiple of 39.1 nanoseconds. Range is 1 to 4294967295.
load	Effective bandwidth of the route expressed as a number from 1 to 255 (255 is 100-percent loaded).
reliability	Likelihood of successful packet transmission expressed as a number between 0 and 255. The value 255 means 100-percent reliability; 0 means no reliability.

Command Default

Metric values are not set.

Command Modes

Interface configuration

Command History

Release	Modification
Release 3.6.2	This command was introduced.

Usage Guidelines

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

Use the **metric** command to provide metric values while redistributing a protocol into an EIGRP interface. Metric defaults have been carefully set to work for a wide variety of networks. Take great care when changing these values.

Task ID

Task ID	Operation
eigrp	read, write

This example shows how to configure metric values for interface POS 0/1/0/1 with values as bandwidth = 100, delay = 7, reliability = 250, and load = 100.

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router eigrp 100
RP/0/RSP0/CPU0:router(config-eigrp)# address-family ipv4
RP/0/RSP0/CPU0:router(config-eigrp-af)# interface GigabitEthernet 0/1/0/1
```

```
RP/0/RSP0/CPU0:router(config-eigrp-af-if)# metric bandwidth 100 delay 7 reliability 250
load 100
```

Related Commands

Command	Description
default-metric (EIGRP), on page 730	Sets metrics for an Enhanced Interior Gateway Routing Protocol (EIGRP).

metric maximum-hops

To advertise as unreachable those Enhanced Interior Gateway Routing Protocol (EIGRP) routes with a hop count higher than is specified by the command, use the **metric maximum-hops** command in the appropriate configuration mode. To reset the value to the default, use the **no** form of this command.

metric maximum-hops *hops-number*
no metric maximum-hops

Syntax Description

hops-number	Maximum hop count. Range is from 1 to 255 hops.

Command Default

hops-number : 100

Command Modes

IPv4 address family configuration

IPv4 VRF address family 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **metric maximum-hops** command to provide a safety mechanism that breaks any potential *count-to-infinity* problems. This command causes the EIGRP routing protocol to advertise as unreachable routes with a hop count greater than the value assigned to the *hops-number* argument.

Task ID

Task ID	Operations
eigrp	read, write

Examples

The following example shows how to configure a hop count to 200 for a router that has a complex WAN generating a large hop count under normal (nonlooping) operations.

```
RP/0/RSP0/CPU0:router(config)# router eigrp 1
RP/0/RSP0/CPU0:router(config-eigrp) address-family ipv4
RP/0/RSP0/CPU0:router(config-eigrp-af)# metric maximum-hops 200
```

Related Commands

Command	Description
metric weights, on page 749	Allows the tuning of the EIGRP metric calculations.

metric rib-scale

To set a RIB scale for EIGRP, use the **metric rib-scale** command in an appropriate configuration mode. To remove the RIB scale and restore the default state, use the **no** form of this command.

metric rib-scale *scale-down-value*
no metric rib-scale

Syntax Description

scale-down-value	Amount to divide the EIGRP Wide Metric by to convert to a 4 Byte RIB metric. Legal ranges are 1-256. Results are expressed as whole integers only (no rounding).

Command Default

Default RIB scale is 128 in the 64 bit mode.

In 32 bit mode, rib scale is always 1.

Command Modes

IPv4 address family configuration

IPv6 address family configuration

IPv4 VRF address family configuration

IPv6 VRF address family configuration (Only supported in 64 bit mode)

Command History

Release	Modification
Release 4.3.0	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.

The ability to scale the RIB metric is provided through the use of the **metric rib-scale** configuration command. When entered, the command will result in all routes in the RIB being cleared and replaced with the new metric values.

Task ID

Task ID	Operation
eigrp	read, write

This example shows how to set the metric rib-scale as 64:

```
Router#configure
Router(config)#router eigrp 1
Router(config-eigrp)#address-family ipv4
Router(config-eigrp-af)#metric rib-scale 64
```

metric version (EIGRP)

By default, EIGRP uses a 64-bit metric to calculate route metrics. To switch to a 32-bit metric version for route metric calculation, use the **metric version** command in the appropriate configuration mode.

metric version { **32bit** }

Syntax Description

32bit	Changes the EIGRP route metric calculation from the default 64-bit version to 32-bit version.

Command Default

By default, EIGRP uses 64-bit metric calculation.

Command Modes

IPv4 address family configuration

Command History

Release	Modification
Release 5.2.0	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.

Use the **metric version** command within the EIGRP address-family configuration mode to switch the metric calculation to 32-bit.

Task ID

Task ID	Operations
eigrp	read, write

Examples

The following example shows how to change the EIGRP route metric calculation to 32-bit version for the IPv4 address family:

```
Router# configure terminal
Router(config)# router eigrp 200
Router(config-eigrp)# address-family ipv4
Router(config-eigrp-af)# metric version 32bit
```

Related Commands

Command	Description
default-metric (EIGRP)	Sets default metric values for redistributed routes

metric weights

To allow the tuning of the Enhanced Interior Gateway Routing Protocol (EIGRP) metric calculations, use the **metric weights** command in the appropriate configuration mode. To reset the values to their defaults, use the **no** form of this command.

metric weights *tos k1 k2 k3 k4 k5*
no metric weights

Syntax Description

tos	Type of service (ToS) which must always be 0.
k1 k2 k3 k4 k5	Constants that convert an EIGRP metric vector into a scalar quantity. The range is 0 to 4294967295.

Command Default

tos: 0

k1: 1

k2: 0

k3: 1

k4: 0

k5: 0

Command Modes

IPv4 address family configuration

IPv4 VRF address family 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **metric weights** command to alter the default behavior of EIGRP routing and metric computation and allow the tuning of the EIGRP metric calculation for a particular ToS.

If k5 equals 0, the composite EIGRP metric is computed according to the following formula:

metric = [k1 * bandwidth + (k2 * bandwidth)/(256 – load) + k3 * delay]

If k5 does not equal zero, an additional operation is performed:

metric = metric * [k5/(reliability + k4)]

Bandwidth is inverse minimum bandwidth of the path in BPS scaled by a factor of 2.56 * 1012. The range is from a 1200-bps line to 10 terabits per second.

Delay is in units of 10 microseconds. The range of delay is from 10 microseconds to 168 seconds. A delay of all 1s indicates that the network is unreachable.

The delay parameter is stored in a 32-bit field, in increments of 39.1 nanoseconds. The range of delay is from 1 (39.1 nanoseconds) to hexadecimal FFFFFFFF (decimal 4,294,967,040 nanoseconds). A delay of all 1s (that is, a delay of hexadecimal FFFFFFFF) indicates that the network is unreachable.

This table lists the default values used for several common media.

Table 50: Bandwidth Values by Media Type

Media Type	Delay	Bandwidth
Satellite	5120 (2 seconds)	5120 (500 megabits)
Ethernet	25600 (1 milliseconds [ms])	256000 (10 megabits)
1.544 Mbps	512000 (20,000 ms)	1,657,856 bits
64 kbps	512000 (20,000 ms)	40,000,000 bits
56 kbps	512000 (20,000 ms)	45,714,176 bits
10 kbps	512000 (20,000 ms)	256,000,000 bits
1 kbps	512000 (20,000 ms)	2,560,000,000 bits

Reliability is given as a fraction of 255. That is, 255 is a reliability of 100 percent or a perfectly stable link. Load is given as a fraction of 255. A load of 255 indicates a completely saturated link.

Task ID

Task ID	Operations
eigrp	read, write

Examples

The following example shows how to set the metric weights to change the default values:

```
RP/0/RSP0/CPU0:router(config)# router eigrp 1
RP/0/RSP0/CPU0:router(config-eigrp) address-family ipv4
RP/0/RSP0/CPU0:router(config-eigrp-af)# metric weights 0 2 0 2 0 0
```

Related Commands

Command	Description
metric maximum-hops, on page 746	Advertises as unreachable those EIGRP VPN routes with a hop count higher than is specified by the command.

neighbor (EIGRP)

To define a neighboring router with which to exchange Enhanced Interior Gateway Routing Protocol (EIGRP) information, use the **neighbor** command in interface configuration mode. To remove an entry, use the **no** form of this command.

neighbor *ip-address interface* [**remote** [*max-hops*]]
noneighbor *ip-address interface* [**remote** [*max-hops*]]

Syntax Description

ip-address	IP address of a peer router with which routing information is exchanged.
interface	Interface through which peering is established.
remote	Specifies that the neighbor is remote.
max-hops	The maximum number of hops within which the neighbor is expected to be reachable from the configured router. The default value is 100 hops.

Command Default

No neighboring routers are defined.

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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **neighbor** command to permit the point-to-point (nonbroadcast) exchange of routing information.

If a neighbor is configured on an interface using the **neighbor** command, the interface stops sending or receiving multicast hello messages. However, the interface can send or receive unicast hello messages. So each neighbor on a LAN must be configured individually .Multiple **neighbor** commands can be used to specify additional neighbors or peers.

Task ID

Task ID	Operations
eigrp	read, write

Examples

This example shows how to permit the sending of EIGRP updates to specific neighbors. One copy of the update is generated for each neighbor:

```
RP/0/RSP0/CPU0:router(config)# router eigrp 100
RP/0/RSP0/CPU0:router(config-eigrp) address-family ipv4
RP/0/RSP0/CPU0:router(config-eigrp-af)# interface GigabitEthernet 0/0/0/3
```

```
RP/0/RSP0/CPU0:router(config-eigrp-af-if)# neighbor 172.20.1.2 remote 10
```

Related Commands

Command	Description
passive-interface (EIGRP), on page 759	Disables sending and receiving "hello" messages on (EIGRP) interface.

neighbor maximum-prefix

To limit the number of prefixes that are accepted from a single Enhanced Interior Gateway Routing Protocol (EIGRP) neighbor or all EIGRP VPN neighbors, use the **neighbor maximum-prefix** command in IPv4 VRF address family configuration mode. To disable this function, use the **no** form of this command.

Single-Neighbor Configuration CLI
neighbor *ip-address* **maximum-prefix** *maximum* [*threshold*] [**warning-only**]
no neighbor *ip-address* **maximum-prefix**

All-Neighbor Configuration CLI
neighbor maximum-prefix *maximum* [*threshold*] [**dampened**] [**reset-time** *minutes*] [**restart** *minutes*] [**restart-count** *number*] [**warning-only**]
no neighbor maximum-prefix

Syntax Description

ip-address	(Optional) IP address of a single peer.
maximum	Maximum number of prefixes accepted. The range is from 1 to 4294967295. The number of prefixes that can be configured is limited only by the available system resources on the router.
threshold	(Optional) Syslog warning messages are specified as a percentage of the maximum prefix limit that was exceeded. The prefix percentage number range is from 1 to 100. The default is 75 percent.
dampened	(Optional) Configures a decay penalty to be applied to the restart time period each time the maximum prefix limit is exceeded. The half-life for the decay penalty is 150 percent of the default or user-defined restart time value in minutes. This keyword is disabled by default.
reset-time *minutes*	(Optional) Configures the router to reset the restart count to 0 after the default or configured reset time period has expired. The range is from 1 to 65535 minutes.
restart *minutes*	(Optional) Configures a time period in which the router does not form adjacencies or accept redistributed routes from the RIB after the maximum prefix limit has been exceeded. The range is from 1 to 65535 minutes.
restart-count *number*	(Optional) Configures the number of times a peering session can be automatically reestablished after the peering session has been torn down or a redistribute route has been cleared and relearned because the maximum prefix limit has been exceeded. **Caution** After the restart count threshold is crossed, you need to use the **clear eigrp neighbors** command to reestablish normal peering, redistribution, or both.
warning-only	(Optional) Configures the router to generate syslog messages only when the maximum prefix limit is reached, instead of terminating the peering session.

Command Default

threshold: 75 percent

dampened : disabled

warning-only : disabled

reset-time : 15 minutes

restart : 5 minutes

restart-count : 3

Command Modes

IPv4 VRF address family 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

The **neighbor maximum-prefix**command is configured to protect an individual peering session or all peering sessions. When this feature is enabled and the maximum prefix limit is exceeded, the router tears down the peering session, clears all routes that are learned from the peer, and then places the peer in a penalty state for the default or user-defined time period. After the penalty time period expires, normal peering is reestablished.

Note In EIGRP, **neighbor** commands have been used traditionally to configure static neighbors. In the context of this feature, however, the **neighbor maximum-prefix** command can be used to configure the maximum prefix limit for both statically configured and dynamically discovered neighbors.

When configuring the **neighbor maximum-prefix** command to protect a single peering session, just the maximum prefix limit, percentage threshold, and warning only configuration options can be configured. Session dampening, restart, and reset timers are configured on a global basis.

Task ID

Task ID	Operations
eigrp	read, write

Examples

The following example shows how to configure the maximum prefix limit for a single peer. The maximum limit is set to 1000 prefixes, and the warning threshold is set to 80 percent. When the maximum prefix limit is exceeded, the session with this peer is torn down, all routes learned from this peer are removed from the topology and routing tables, and this peer is placed in a penalty state for 5 minutes (default penalty value).

```
RP/0/RSP0/CPU0:router(config)# router eigrp 1
RP/0/RSP0/CPU0:router(config-eigrp)# vrf vrf1
RP/0/RSP0/CPU0:router(config-eigrp-vrf)# address-family ipv4
RP/0/RSP0/CPU0:router(config-eigrp-vrf-af)# neighbor 10.0.0.1 maximum-prefix 1000 80
```

The following example shows how to configure the maximum prefix limit for all peers. The maximum limit is set to 10,000 prefixes, the warning threshold is set to 90 percent, the restart timer is set to 4 minutes, a decay penalty is configured for the restart timer with the **dampened** keyword, and all timers are configured to be reset to 0 every 60 minutes. When the maximum prefix limit is exceeded, all peering sessions are torn down, all routes learned from all peers are removed from the topology and routing tables, and all peers are placed in a penalty state for 4 minutes (user-defined penalty value). A dampening exponential decay penalty is also applied.

```
  RP/0/RSP0/CPU0:router(config)# router eigrp 1
  RP/0/RSP0/CPU0:router(config-eigrp)# vrf vrf1
  RP/0/RSP0/CPU0:router(config-eigrp-vrf)# address-family ipv4
  RP/0/RSP0/CPU0:router(config-eigrp-vrf-af)# neighbor maximum-prefix 10000 90 dampened
reset-time 60 restart4
```

Related Commands

Command	Description
clear eigrp neighbors, on page 725	Deletes EIGRP and VRF neighbor entries from the appropriate tables.

next-hop-self disable

To instruct the Enhanced Interior Gateway Routing Protocol (EIGRP) process to use the received next-hop value when advertising the routes, use the **next-hop-self disable**command in interface configuration mode. To revert to the default, use the **no** form of this command.

next-hop-self disable
no next-hop-self disable

Syntax Description

This command has no keywords or arguments.

Command Default

EIGRP always sets the IP next-hop value to be itself.

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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

EIGRP, by default, sets the IP next-hop value to be itself for routes that it is advertising, even when advertising those routes on the same interface from which learned them. To change this default, you must use the **next-hop-self disable** interface configuration command to instruct EIGRP to use the received next-hop value when advertising these routes.

The **next-hop-self disable** feature is not available for redistributed routes.

Task ID

Task ID	Operations
eigrp	read, write

Examples

The following example shows how to change the default IP next-hop value and instruct EIGRP to use the received next-hop value:

```
RP/0/RSP0/CPU0:router(config)# router eigrp 1
RP/0/RSP0/CPU0:router(config-eigrp) address-family ipv4
RP/0/RSP0/CPU0:router(config-eigrp-af)# interface GigabitEthernet 0/1/0/0
RP/0/RSP0/CPU0:router(config-eigrp-af-if)# next-hop-self disable
```

nsf disable (EIGRP)

To disable Enhanced Interior Gateway Routing Protocol (EIGRP) nonstop forwarding (NSF), use the **nsf disable** command in appropriate configuration mode. To re-enable the EIGRP NSF from a disabled state, use the **no** form of this command.

nsf disable
no nsf disable

Syntax Description

This command has no keywords or arguments.

Command Default

NSF in EIGRP is enabled.

Command Modes

Router Configration

IPv4 address family configuration

IPv6 address family configuration

IPv4 VRF address family configuration

IPv6 VRF address family configuration

Command History

Release	Modification
Release 4.1.1	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.

If NSF is to be disabled for both IPv4 and IPv6 address families of all VRFs, use the **nsf disable** command at router configuration mode.

If NSF is to be disabled for a specific address family under a specific VRF, use the **nsf disable** command at address family configuration mode.

If NSF is disabled, EIGRP will not be able to retain the routes learnt from its peers and may result in traffic loss during ISSU.

Task ID

Task ID	Operation
eigrp	read, write

This example shows how to disable NSF for all address families under all VRF's:

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#router eigrp 1
```

```
RP/0/RSP0/CPU0:router(config-eigrp)#nsf disable
```

This example shows how to disable NSF for IPv4 address family of VRF v1:

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#router eigrp 1
RP/0/RSP0/CPU0:router(config-eigrp)#vrf v1
RP/0/RSP0/CPU0:router(config-eigrp-vrf)#address-family ipv4
RP/0/RSP0/CPU0:router(config-eigrp-vrf-af)#nsf disable
```

Related Commands

Command	Description
router eigrp, on page 767	Configures a routing process and enters router configuration mode for Enhanced Interior Gateway Routing Protocol (EIGRP).

passive-interface (EIGRP)

To disable sending and receiving "hello" messages on an Enhanced Interior Gateway Routing Protocol (EIGRP) interface and to disable formation of neighbors on the interface, use the **passive-interface** command in interface configuration mode. To reenable sending and receiving "hello messages", use the **no** form of this command.

passive-interface
no passive-interface

Syntax Description

This command has no keywords or arguments.

Command Default

passive-interface command is disabled on an interface.

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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **passive-interface** command to disable sending of "hello" massages. The particular subnet on that interface will continue to be advertised by EIGRP to neighbors on other interfaces.

Task ID

Task ID	Operations
eigrp	read, write

Examples

This example shows how to configure **passive-interface** command on GigabitEthernet interface 0/6/5/0:

```
RP/0/RSP0/CPU0:router(config)# router eigrp 1
RP/0/RSP0/CPU0:router(config-eigrp) address-family ipv6
RP/0/RSP0/CPU0:router(config-eigrp-af)# interface GigabitEthernet 0/6/5/0
RP/0/RSP0/CPU0:router(config-eigrp-af-if)# passive-interface
```

redistribute (EIGRP)

To inject routes from one routing domain into the Enhanced Interior Gateway Routing Protocol (EIGRP), use the **redistribute** command in the appropriate configuration mode. To remove the **redistribute** command from the configuration file and restore the system to its default condition in which the software does not redistribute routes, use the **no** form of this command.

redistribute {**{bgp** | **connected** | **isis** | **ospf** | **rip** | **static** | **eigrp}** | [*as-numberinstance-name*]}[**route-policy** *name*]
no redistribute

Syntax Description

bgp	Distributes routes from the BGP protocol.
connected	Distributes routes that are established automatically by virtue of having enabled IP on an interface.
isis	Distributes routes from the IS-IS protocol.
ospf	Distributes routes from the OSPF protocol. This protocol is supported in the IPv4 address family.
static	Redistributes IP static routes.
as-number | *instance-name*	Represents one of the following three options: For the **bgp** keyword: 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. For the **isis** keyword, an IS-IS instance name from which routes are to be redistributed. The value takes the form of a string. A decimal number can be entered, but it is stored internally as a string. For the **ospf** keyword, an OSPF instance name from which routes are to be redistributed. The value takes the form of a string. A decimal number can be entered, but it is stored internally as a string.
route-policy *name*	(Optional) Specifies the identifier of a configured policy. A policy is used to filter the importation of routes from this source routing protocol to EIGRP.

Command Default

Route redistribution is disabled.

Command Modes

IPv4 address family configuration

IPv4 VRF address family configuration

Command History

Release	Modification
Release 3.7.2	This command was introduced.
Release 3.9.0	Asplain format for 4-byte Autonomous system numbers notation was supported.

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.

Redistributed routing information should always be filtered by the **route- policy** *name* keyword and argument. This filtering ensures that only those routes intended by the administrator are redistributed by EIGRP.

A default metric is usually required to redistribute routes from another protocol into EIGRP. The metric is configured through the **default-metric** command or under the route policy configured with the **redistribute** command. The one exception to this requirement is when EIGRP redistributes BGP routes on a provider edge (PE) router in an MPLS-VPN scenario. If the originating protocol of the route is EIGRP with the same autonomous system (as in MPLS-VPN), the metric would be learned automatically from the extended communities of the BGP route.

For information about routing policies, see the *Routing Policy Commands on* Cisco ASR 9000 Series Router module of the *Cisco ASR 9000 Series Aggregation Services Router Routing Command Reference*.

Task ID

Task ID	Operations
eigrp	read, write

Examples

This example shows how to cause BGP routes to be redistributed into an EIGRP autonomous system:

```
RP/0/RSP0/CPU0:router(config)# router eigrp 1
RP/0/RSP0/CPU0:router(config-eigrp) address-family ipv4
RP/0/RSP0/CPU0:router(config-eigrp-af)# redistribute bgp 100
```

This example shows how to redistribute the specified IS-IS process routes into an EIGRP autonomous system within a VPN routing and forwarding instance (VRF). The IS-IS routes are redistributed using route policy 3.

```
RP/0/RSP0/CPU0:router(config)# router eigrp 109
RP/0/RSP0/CPU0:router(config-eigrp)# vrf vpn-1
RP/0/RSP0/CPU0:router(config-eigrp-vrf)# address-family ipv4
RP/0/RSP0/CPU0:router(config-eigrp-vrf-af)# redistribute isis 108 route-policy 3
```

Related Commands

Command	Description
default-metric (EIGRP), on page 730	Sets metrics for an EIGRP.

redistribute maximum-prefix

To limit the number of prefixes redistributed into an Enhanced Interior Gateway Routing Protocol (EIGRP) process, use the **redistribute maximum-prefix** command in IPv4 VRF address family configuration mode. To disable this function, use the **no** form of this command.

redistribute maximum-prefix *maximum* [*threshold*] [[**dampened**] [**reset-time** *minutes*] [**restart** *minutes*] [**restart-count** *number*] | [**warning-only**]]
no redistribute maximum-prefix

Syntax Description

maximum	Maximum number of prefixes that are redistributed into EIGRP under an address-family. The range is from 1 to 4294967295. The number of prefixes that can be configured is limited only by the available system resources on the router.
threshold	(Optional) Syslog warning messages are specified as a percentage of the maximum prefix limit that was exceeded. The prefix percentage number range is from 1 to 100. The default is 75 percent.
restart *minutes*	(Optional) Configures a time period in which the router will not form adjacencies or accept redistributed routes from the RIB after the *maximum*-prefix limit has been exceeded. The value for the minutes argument is from 1 to 65535 minutes.
restart-count *number*	(Optional) Configures the number of times a peering session can be automatically reestablished after the peering session has been torn down or after the redistribute route has been cleared and relearned because the maximum prefix limit has been exceeded. After the restart count threshold has been crossed, you will need to enter the **process restart eigrp** command to reestablish normal peering, redistribution, or both.
reset-time *minutes*	(Optional) Configures the router to reset the restart count to 0 after the default or configured reset time period has expired. The value for the minutes argument is from 1 to 65535 minutes.
dampened	(Optional) Configures a decay penalty to be applied to the restart time period each time the maximum prefix limit is exceeded. The half-life for the decay penalty is 150 percent of the default or user-defined restart time value in minutes.
warning-only	(Optional) Configures the router to only generate syslog messages when the maximum prefix limit is reached, instead of suspending redistribution.

Command Default

threshold: 75 percent

warning-only : disabled

reset-time : 15 minutes

restart : 5 minutes

restart-count : 3

dampened: disabled

Command Modes

IPv4 VRF address family 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **redistribute maximum-prefix** command to configure limit prefixes learned through redistribution. When the maximum prefix limit is exceeded, all routes learned from the Routing Information Base (RIB) are discarded and redistribution is suspended for the default or user-defined time period. The maximum prefix limit that can be configured for redistributed prefixes is limited only by the available system resources on the router.

Task ID

Task ID	Operations
eigrp	read, write

Examples

The following example shows how to configure the maximum prefix limit for routes learned through redistribution. The maximum limit is set to 5000 prefixes, and the warning threshold is set to 95 percent. When the number of prefixes learned through redistribution reaches 4750 (95 percent of 5000), warning messages are displayed in the console. Because the **warning-only** keyword is configured, the topology and routing tables are not cleared and route redistribution is not placed in a penalty state.

```
RP/0/RSP0/CPU0:router(config)# router eigrp 100
RP/0/RSP0/CPU0:router(config-eigrp)# vrf vpn-1
RP/0/RSP0/CPU0:router(config-eigrp-vrf)# address-family ipv4
RP/0/RSP0/CPU0:router(config-eigrp-vrf-af)# redistribute maximum-prefix 5000 95 warning-only
```

Related Commands

Command	Description
process	To start, terminate, or restart a process.

remote-neighbor (unicast)

To configure any neighbor that is multiple hops away, including remote static neighbors, use the **remote-neighbor** command.

remote-neighbors unicast-listen [**allow-list** *route-policyname*] [**max-neighbors** *maxRemotePeers*]

no remote-neighbors unicast-listen [**allow-list** *route-policyname*] [**max-neighbors** *maxRemotePeers*]

Syntax Description

unicast-listen	Use unicast to form remote neighbor relationship without having to manually configure the remote neighbors addresses.
allow-list*route policy name*	Name of the route-policy that specifies remote addresses from which EIGRP neighbor connections may be accepted.
max-neighbors*maximum remote peers*	The maximum number of remote neighbors from which connection can be accepted. The range is 1-65535.

Command Modes

Router configuration

Command History

Release	Modification
Release 5.2.2	This command was introduced.

Usage Guidelines

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

Use the **remote-neighbor unicast-listen** command to configure a neighbor to use unicast to communicate with a neighbor that is multiple hops away, and whose address has not been configured with the **neighbor** command. A single unicast address can only be configured to a single remote static neighbor for a given address-family. If you try to configure a second remote static neighbor using the same unicast address but a different interface, it is rejected. EIGRP configuration of remote neighbors under different address families is unrestricted.

Task ID

Task ID	Operation
eigrp	read, write

Example

This example shows you how to configure both devices (hub and spoke) involved in the neighbor relationship.

```
RP/0/RSP0/CPU0:HUB(config)#router eigrp 100
RP/0/RSP0/CPU0:HUB(config-eigrp)#address-family ipv4
RP/0/RSP0/CPU0:HUB(config-eigrp-af)#int g0/0/0/3
```

```
RP/0/RSP0/CPU0:HUB(config-eigrp-af-if)#exit
RP/0/RSP0/CPU0:HUB(config-eigrp-af)#interface gigabitEthernet 0/0/0/3
RP/0/RSP0/CPU0:HUB(config-eigrp-af-if)#remote-neighbor unicast-listen
RP/0/RSP0/CPU0:HUB(config-eigrp-af-if)#commit

RP/0/RSP0/CPU0:spoke(config)#router eigrp 100
RP/0/RSP0/CPU0:spoke(config-eigrp)#address-family ipv4
RP/0/RSP0/CPU0:spoke(config-eigrp-af)#interface g0/0/0/3
RP/0/RSP0/CPU0:spoke(config-eigrp-af-if)#neighbor 21.21.21.1
RP/0/RSP0/CPU0:spoke(config-eigrp-af-if)#commit

RP/0/RSP0/CPU0:spoke#sh run router eigrp
Fri Aug  8 08:47:48.556 UTC
router eigrp 100
address-family ipv4
interface GigabitEthernet0/0/0/3
neighbor 21.21.21.1  !!!
```

route-policy (EIGRP)

To apply a routing policy to updates advertised to or received from an Enhanced Interior Gateway Routing Protocol (EIGRP) neighbor, use the **route-policy** command in the appropriate configuration mode. To disable applying routing policy to updates, use the **no** form of this command.

route-policy *route-policy-name* {**in** | **out**}
no route-policy *route-policy-name* {**in** | **out**}

Syntax Description

route-policy-name	Name of route policy.
in	Applies policy to inbound routes.
out	Applies policy to outbound routes.

Command Default

No policy is applied.

Command Modes

IPv4 address family configuration

IPv4 VRF address family configuration

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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

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.

Task ID

Task ID	Operations
eigrp	read, write

Examples

The following example shows how to apply the IN-Ipv4 policy to inbound IP Version 4 (IPv4) routes:

```
RP/0/RSP0/CPU0:router(config)# router eigrp 1
RP/0/RSP0/CPU0:router(config-eigrp) address-family ipv4
RP/0/RSP0/CPU0:router(config-eigrp-af)# route-policy IN-IPv4 in
```

router eigrp

To configure a routing process and enter router configuration mode for Enhanced Interior Gateway Routing Protocol (EIGRP), use the **router eigrp** command in global configuration mode. To turn off the EIGRP routing process, use the **no** form of this command.

router eigrp { *instance-autonomous-system-number virtual-instance-name* }
no router eigrp { *instance-autonomous-system-number virtual-instance-name* }

Syntax Description

instance-autonomous-system-number	EIGRP instance autonomous-system number. This is used as the autonomous-system number for the default/global VRF. Valid range is 1 to 65535.

Command Default

No routing process is defined.

Command Modes

Global 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

- Only one instance of EIGRP is supported.
-

Task ID

Task ID	Operations
eigrp	read, write

Examples

The following example configures EIGRP routing process 109 (109 is the autonomous-system number of default/global VRF):

```
RP/0/RSP0/CPU0:router(config)# router eigrp 109
RP/0/RSP0/CPU0:router(config-eigrp)#
```

router-id (EIGRP)

To configure a router ID for an Enhanced Interior Gateway Routing Protocol (EIGRP) process, use the **router-id** command in the appropriate configuration mode. To cause the software to use the default method of determining the router ID, use the **no** form of this command.

router-id *router-id*
no router-id

Syntax Description

router-id	32-bit router ID value specified in four-part, dotted-decimal notation.

Command Default

If this command is not configured, EIGRP chooses an IPv4 address as the router ID from one of its interfaces.

Command Modes

IPv4 address family configuration

IPv4 VRF address family 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

It is good practice to use the **router-id** command to explicitly specify a unique 32-bit numeric value for the router ID. This action ensures that EIGRP can function regardless of the interface address configuration.

Task ID

Task ID	Operations
eigrp	read, write

Examples

The following example shows how to assign the IP address of 172.20.1.1 to the EIGRP process 1:

```
RP/0/RSP0/CPU0:router(config)# router eigrp 1
RP/0/RSP0/CPU0:router(config-eigrp) address-family ipv4
RP/0/RSP0/CPU0:router(config-eigrp-af)# router-id 172.20.1.1
```

show eigrp accounting

To display prefix accounting information for Enhanced Interior Gateway Routing Protocol (EIGRP) processes, use the **show eigrp accounting** command in EXEC mode.

show eigrp [*as-number*] [**vrf** {*vrf-name* | **all**}] [**ipv4** | **ipv6**] **accounting**

Syntax Description

as-number	(Optional) Autonomous system number. This option is available when a VPN routing and forwarding (VRF) instance is not specified. Range is from 1 to 65535.
vrf { *vrf-name* | **all** }	(Optional) Specifies a particular VPN routing and forwarding instance (VRF) or all VRF instances.
[**ipv4**]	(Optional) Specifies the IPv4 address family.

Command Default

This command has no arguments or keywords.

Command Modes

EXEC

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 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
EIGRP	read

Examples

The following is sample output from the **show eigrp accounting** command:

```
RP/0/RSP0/CPU0:router# show eigrp accounting

IP-EIGRP accounting for AS(100)/ID(10.0.2.1) Routing Table: RED
Total Prefix Count: 4  States: A-Adjacency, P-Pending, D-Down

State Address/Source   Interface        Prefix   Restart  Restart/
                                        Count     Count   Reset(s)
 P    Redistributed     ----               0         3         211
 A    10.0.1.2          Gi0/6/0/0/         2         0          84
 P    10.0.2.4          Gi00/2/0/3         0         2         114
 D    10.0.1.3          Gi0/6/0/0          0         3           0
```

Note Connected and summary routes are not listed individually in the output of this command but are counted in the total aggregate count for each process.

This table describes the significant fields shown in the display.

Table 51: show eigrp accounting Field Descriptions

Field	Description
EIGRP accounting for AS	Identifies the EIGRP instance along with the AS number, router ID and table ID.
Total Prefix Count	Shows the aggregate sum of the prefixes in an EIGRP instance topology table. The count includes prefixes learned from all neighbors or from redistribution.
States: A-Adjacency, P-Pending, D-Down	A-Adjacency: Indicates a stable adjacency with the neighbor or a normal redistribution state. P-Pending: Neighbor adjacency or redistribution is suspended or in a penalized state because the maximum prefix limit was exceeded. D-Down: Neighbor adjacency or redistribution is suspended permanently until a manual reset is performed with the **clear route** command.
Address/Source	Shows the peer IP address of the redistribution source.
Prefix Count	Displays the total number of learned prefixes by source. **Note** Routes can be learned for the same prefix from multiple sources, and the sum of all prefix counts in this column may be greater than the figure displayed in the "Prefix Count" field.
Restart Count	Number of times a route source exceeded the maximum prefix limit.
Restart Reset(s)	Displays the time, in seconds, that a route source is in a P (penalized) state. If the route source is in an A (stable or normal) state, the displayed time, in seconds, is the time period until penalization history is reset.

show eigrp interfaces

To display information about interfaces configured for Enhanced Interior Gateway Routing Protocol (EIGRP), use the **show eigrp interfaces** command in EXEC mode.

show eigrp [*as-number*] [**vrf** {*vrf-name* | **all**}] [**ipv4** | **ipv6**] **interfaces** [*type interface-path-id*] [**detail**]

Syntax Description

as-number	(Optional) Autonomous system number. This option is available when a VPN routing and forwarding (VRF) instance is not specified. 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.
vrf { *vrf-name* | **all** }	(Optional) Specifies a particular VPN routing and forwarding instance (VRF) or all VRF instances.
[**ipv4**]	(Optional) Specifies the IPv4 address family.
type	(Optional) Interface type. For more information, use the question mark (**?**) online help function.
interface-path-id	Physical interface or virtual interface. 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.
detail	(Optional) Displays detailed EIGRP interface information.

Command Default

This command has no arguments or keywords.

Command Modes

EXEC

Command History

Release	Modification
Release 3.7.2	This command was introduced.
Release 3.9.0	As plain format for 4-byte Autonomous system numbers notation was supported. The input parameters and output were modified to display 4-byte autonomous system numbers and extended communities in either asplain or asdot notations.

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.

Use the **show eigrp interfaces** command to determine on which interfaces EIGRP is active and learn information about EIGRP related to those interfaces.

If an interface is specified, only that interface is displayed. Otherwise, all interfaces on which EIGRP is running are displayed.

If an autonomous system is specified, only the routing process for the specified autonomous system is displayed. Otherwise, all EIGRP processes are displayed.

Task ID

Task ID	Operations
EIGRP	read

Examples

The following is sample output from the **show eigrp interfaces** command:

```
RP/0/RSP0/CPU0:router# show eigrp interfaces

IP EIGRP interfaces for process 1

                        Xmit Queue    Mean   Pacing Time   Multicast   Pending
Interface        Peers  Un/Reliable   SRTT   Un/Reliable   Flow Timer  Routes

Gi0/6/0/2.212      0         0/0          0      11/434          0          0

Gi0/6/0/0          1         0/0        337       0/10           0          0

 Gi0/2/0/3         1         0/0         10       1/63         103          0

Gi0/6/2/5          1         0/0        330       0/16           0          0
```

This table describes the significant fields shown in the display.

Table 52: show eigrp interfaces Field Descriptions

Field	Description
Interface	Interface over which EIGRP is configured.
Peers	Number of directly connected EIGRP neighbors.
Xmit Queue Un/Reliable	Number of packets remaining in the unreliable and reliable transmit queues.
Mean SRTT	Mean smoothed round-trip time (SRTT) internal (in milliseconds).
Pacing Time Un/Reliable	Pacing time used to determine when EIGRP packets should be sent out the interface (unreliable and reliable packets).
Multicast Flow Timer	Maximum number of seconds in which the router sends multicast EIGRP packets.
Pending Routes	Number of routes in the packets in the transmit queue waiting to be sent.

The following is sample output from the **show eigrp interfaces** command when issued with the **detail** keyword:

```
RP/0/RSP0/CPU0:router# show eigrp interfaces detail

IPv4-EIGRP interfaces for AS(100)

                        Xmit Queue   Mean   Pacing Time   Multicast    Pending
Interface        Peers  Un/Reliable  SRTT   Un/Reliable   Flow Timer   Routes
Lo0                0        0/0         0     640/640          0           0
  Hello interval is 5 sec, hold time is 15 sec
  Next xmit serial <none>
  Un/reliable mcasts: 0/0  Un/reliable ucasts: 0/0
  Mcast exceptions: 0  CR packets: 0  ACKs suppressed: 0
  Retransmissions sent: 0  Out-of-sequence rcvd: 0
  Bandwidth percent is 50
  Total packets received: 0
  Authentication mode: MD5  Key chain: key1
  No active key found in keychain database
  Valid authenticated packets received: 0
  Packets dropped due to wrong keychain config: 0
  Packets dropped due to missing authentication: 0
  Packets dropped due to invalid authentication: 0
  Effective Metric:
    Bandwidth: 10000000, Delay: 500, Reliability: 255, Load: 1, MTU: 1514
```

This table describes the significant fields shown in the display.

Table 53: show eigrp interfaces detail Field Descriptions

Field	Description
Hello interval	Hello packet transmission interval.
hold time	Hold time announced to neighbors. If neighbors do not get hello packets from the router for this period, neighbors declare that the neighbor relationship is down.
Next xmit serial	Next transmission serial number.
Un/reliable mcasts	Number of unreliable and reliable multicast packets sent on this interface.
Un/reliable ucasts	Number or unreliable and reliable unicast packets sent on this interface.
Mcast exceptions	Number of multicast exceptions (sequence TLVs).
CR packets	Number of packets sent with the conditional receive bit set.
ACKs suppresses	Number of ACK packets suppressed.
Retransmissions	Number of retransmissions sent on this interface.
Out-of-sequence rcvd	Number of packets received out of sequence.
Bandwidth percent	Configured percent of bandwidth.
Authentication	Mode of authentication.

Field	Description
Valid authenticated packets received	Number of valid authentication packets.
Packets dropped due to wrong keychain config	Number of packets dropped due to wrong keychain configuration.
Packets dropped due to missing authentication	Number of packets dropped due to missing authentication.
Packets dropped due to invalid authentication	Number of packets dropped due to invalid authentication.

Related Commands

Command	Description
show eigrp neighbors, on page 775	Displays the neighbors discovered by EIGRP.

show eigrp neighbors

To display information about neighbors discovered by Enhanced Interior Gateway Routing Protocol (EIGRP), use the **show eigrp neighbors** command in EXEC mode.

show eigrp *as-numbervrf* {*vrf-name* | **all**}**ipv4ipv6**

Syntax Description

as-number	(Optional) Autonomous system number. This option is available when a VPN routing and forwarding (VRF) instance is not specified. Range is from 1 to 65535.
vrf { *vrf-name* | **all** }	(Optional) Specifies a particular VPN routing and forwarding instance (VRF) or all VRF instances.
[**ipv4**]	(Optional) Specifies the IPv4 address family.
detail	(Optional) Displays detailed EIGRP neighbor information.
type	Interface type. For more information, use the question mark (**?**) online help function.
interface-path-id	Physical interface or virtual interface. 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.
static	(Optional) Displays static routes.

Command Default

No default behavior or values

Command Modes

EXEC

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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **show eigrp neighbors** command to determine when neighbors become active and inactive. This command is also useful for debugging certain types of transport problems.

Task ID

Task ID	Operations
EIGRP	read

Examples

The following is sample output from the **show eigrp neighbors** command:

```
RP/0/RSP0/CPU0:router# show eigrp neighbors

IP-EIGRP Neighbors for process 77

Address                 Interface     Holdtime Uptime   Q      Seq  SRTT  RTO
                                      (secs)   (h:m:s)  Count  Num  (ms)  (ms)

172.16.81.28             Gi0/2/0/3     13       0:00:41  0      11   4     20

172.16.80.28            Gi0/6/0/0      14       0:02:01  0      10   12    24

172.16.80.31            Gi0/6/2/5      12       0:02:02  0      4    5     20

RP/0/RSP0/CPU0:router#
```

This table describes the significant fields shown in the display.

Table 54: show eigrp neighbors Field Descriptions

Field	Description
process	Autonomous system number specified in the router configuration command.
Address	IP address of the EIGRP peer.
Interface	Interface on which the router is receiving hello packets from the peer.
Holdtime	Length of time (in seconds) that the Cisco IOS XR software waits to hear from the peer before declaring that the peer is down.
Uptime	Elapsed time (in hours, minutes, and seconds) since the local router first heard from this neighbor.
Q Count	Number of EIGRP packets (update, query, and reply) that the software waits to send.
Seq Num	Sequence number of the last update, query, or reply packet that was received from this neighbor.
SRTT	Smoothed round-trip time. This is the number of milliseconds required for an EIGRP packet to be sent to this neighbor and for the local router to receive an acknowledgment of that packet.
RTO	Retransmission timeout (in milliseconds). This is the amount of time the software waits before resending a packet from the retransmission queue to a neighbor.

The following is sample output from the **show eigrp neighbors** command when issued with the **detail** keyword:

```
RP/0/RSP0/CPU0:router# show eigrp neighbors detail

IP-EIGRP neighbors for AS 1

H   Address                 Interface       Hold Uptime   SRTT   RTO  Q  Seq
                                            (sec)         (ms)       Cnt Num
0   11.0.0.10               Gi0/6/0/0         14 01:00:52    3   200  0  10

   Version 12.4/1.2, Retrans: 0, Retries: 0, Prefixes: 3
```

This table describes the significant fields shown in the display.

Table 55: show eigrp neighbors detail Field Descriptions

Field	Description
Version	Version of EIGRP software (major.minor) running on the node and neighbor.
Retrans	Number of retransmissions sent to this neighbor.
Retries	Number of retransmissions sent to this neighbor since the last acknowledgement (ACK).
Prefixes	Number of prefixes learned from this neighbor.

show eigrp topology

To display the Enhanced Interior Gateway Routing Protocol (EIGRP) topology table, use the **show eigrp topology** command in EXEC mode.

show eigrp [*as-number*] [**vrf** {*vrf-name* | **all**}] [**ipv4** | **ipv6**] **topology** [*ip-address mask*] {**active** | **all-links** | **detail-links** | **pending** | **summary** | **zero-successors**}

Syntax Description

as-number	(Optional) Autonomous system number. This option is available when a VPN routing and forwarding (VRF) instance is not specified. 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.
vrf { *vrf-name* | **all** }	(Optional) Specifies a particular VPN routing and forwarding instance (VRF) or all VRF instances.
[**ipv4**]	(Optional) Specifies the IPv4 address family.
ip-address	(Optional) IP address in four-part, dotted-decimal notation.
mask	(Optional) Network mask specified in either of two ways: Network mask can be a four-part, dotted decimal address. For example, 255.0.0.0 indicates that each bit equal to 1 means the corresponding address bit is a network address. Network mask can be indicated as a slash (/) and number. For example, /8 indicates that the first 8 bits of the mask are 1s, and the corresponding bits of the address are the network address.
active	(Optional) Displays only active entries in the EIGRP topology table.
all-links	(Optional) Displays all entries in the EIGRP topology table.
detail-links	(Optional) Displays detailed information for all entries in the EIGRP topology table.
pending	(Optional) Displays all entries in the EIGRP topology table that are waiting for an update from a neighbor or are waiting to reply to a neighbor.
summary	(Optional) Displays a summary of the EIGRP topology table.
zero-successors	(Optional) Displays available routes in the EIGRP topology table.

Command Modes

EXEC

Command History

Release	Modification
Release 3.7.2	This command was introduced.
Release 3.9.0	Asplain format for 4-byte Autonomous system numbers notation was supported. The input parameters and output were modified to display 4-byte autonomous system numbers and extended communities in either asplain or asdot notations.

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.

When the **show eigrp topology** command is used without any keywords or arguments, only routes that are feasible successors are displayed.

The **show eigrp topology** command can be used to determine Diffusing Update Algorithm (DUAL) states and to debug possible DUAL problems.

Task ID

Task ID	Operations
EIGRP	read

Examples

The following is sample output from the **show eigrp topology** command. The EIGRP metrics for specified internal routes and external routes are displayed.

```
RP/0/RSP0/CPU0:router# show eigrp topology 10.2.1.0/24

IP-EIGRP (AS 1): Topology entry for 10.2.1.0/24
  State is Passive, Query origin flag is 1, 1 Successor(s), FD is 281600
  Routing Descriptor Blocks:
  0.0.0.0 (GigabitEthernet0/6/0/0), from Connected, Send flag is 0x0
    Composite metric is (281600/0), Route is Internal !This is the internal route.
    Vector metric:
      Minimum bandwidth is 10000 Kbit
      Total delay is 1000 microseconds
      Reliability is 255/255
      Load is 1/255
      Minimum MTU is 1500
      Hop count is 0

RP/0/RSP0/CPU0:router# show eigrp topology 10.4.80.0/20

IP-EIGRP (AS 1): Topology entry for 10.4.80.0/20
  State is Passive, Query origin flag is 1, 1 Successor(s), FD is 409600
  Routing Descriptor Blocks:
  10.2.1.1 (GigabitEthernet0/6/0/0), from 10.2.1.1, Send flag is 0x0
    Composite metric is (409600/128256), Route is External
    Vector metric:
      Minimum bandwidth is 10000 Kbit
      Total delay is 6000 microseconds
      Reliability is 255/255
      Load is 1/255
      Minimum MTU is 1500
      Hop count is 1
    External data:
      Originating router is 10.89.245.1
```

```
AS number of route is 0
External protocol is Connected, external metric is 0
Administrator tag is 0 (0x00000000)
```

This table describes the significant fields shown in the display.

Table 56: show eigrp topology Field Descriptions

Field	Description
Query origin	Query origin state.
Successors	Number of feasible successors for this prefix.
FD	Feasible distance for this prefix.
10.2.1.1 (Gi0/0)	Next hop and interface from which this path was learned.
from 10.2.1.1	Information source for this path.
Send flag	Indicates if the sending of this prefix is pending to this neighbor.
Composite Metric (409600/128256)	The first number is the EIGRP metric that represents the cost to the destination. The second number is the EIGRP metric that this peer advertised.
Route is	Type of route (internal or external).
Vector Metric	Shows the metric (bandwidth, delay, reliability, load, MTU, and hop count) advertised by the neighbor.
External Data	Shows the external information (originating router ID, AS number, external protocol, metric, and tag) advertised by the neighbor.

show eigrp traffic

To display the number of Enhanced Interior Gateway Routing Protocol (EIGRP) packets sent and received, use the **show eigrp traffic** command in EXEC mode.

show eigrp [*as-number*] [**vrf** {*vrf-name* | **all**}][**ipv4** | **ipv6**] **traffic**

Syntax Description

as-number	(Optional) Autonomous system number. This option is available when a VPN routing and forwarding (VRF) instance is not specified. Range is from 1 to 65535.
vrf { *vrf-name* | **all** }	(Optional) Specifies a particular VPN routing and forwarding instance (VRF) or all VRF instances.
[**ipv4**]	(Optional) Specifies the IPv4 address family.

Command Modes

EXEC

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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **show eigrp traffic** command to find the number of packets sent and received.

In addition, this command is useful in determining whether packets from one node are not reaching the neighboring node due to connectivity or configuration problems.

Task ID

Task ID	Operations
EIGRP	read

Examples

The following is sample output from the **show eigrp traffic** command:

```
RP/0/RSP0/CPU0:router# show eigrp traffic

IP-EIGRP Traffic Statistics for AS 1

  Hellos sent/received: 736/797
  Updates sent/received: 6/6
  Queries sent/received: 0/1
  Replies sent/received: 1/0
  Acks sent/received: 6/6
  Input queue high water mark 0, 0 drops
  SIA-Queries sent/received: 0/0
  SIA-Replies sent/received: 0/0
```

This table describes the significant fields shown in the display.

Table 57: show eigrp traffic Field Descriptions

Field	Description
AS	Autonomous system number specified in the **router eigrp** command.
Hellos sent/received	Number of hello packets sent and received.
Updates sent/received	Number of update packets sent and received.
Queries sent/received	Number of query packets sent and received.
Replies sent/received	Number of reply packets sent and received.
Acks sent/received	Number of acknowledgment packets sent and received.
Input queue high water mark	Maximum number of packets in the input queue and number of drops.
SIA-Queries sent/received	Number of Stuck-in-Active query packets sent and received.
SIA-Replies sent/received	Number of Stuck-in-Active reply packets sent and received.

show protocols (EIGRP)

To display information about the Enhanced Interior Gateway Routing Protocol (EIGRP) process configuration, use the **show protocols** command in EXEC mode.

show protocols [**ipv4** | **afi-all**] [**all** *protocol*] [**default-context** | **vrf***vrf-name*] [**private**]

Syntax Description

ipv4	(Optional) Specifies an IPv4 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 **eigrp** , **bgp** , **isis** , **ospf** , and **rip** .
default-context	(Optional) Displays default context information. This keyword is available when the **eigrp** or **rip** protocol is specified.
vrf *vrf-name*	(Optional) Displays VPN routing and forwarding (VRF) information for the specified process. This keyword is available when the **eigrp** or **rip** protocol is specified.
private	(Optional) Displays private EIGRP data. This keyword is available when the **eigrp** is specified.

Command Modes

EXEC

Command History

Release	Modification
Release 3.7.2	This command was introduced.
Release 3.9.0	Asplain format for 4-byte Autonomous system numbers notation was supported. The input parameters and output were modified to display 4-byte autonomous system numbers and extended communities in either asplain or asdot notations.

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.

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 EIGRP, the command output lists the instance number, default AS context, router ID, default networks, distance, maximum paths, and so on.

Task ID

Task ID	Operations
EIGRP	read

Examples

The following is sample output from the **show protocols eigrp** command:

```
RP/0/RSP0/CPU0:router# show protocols eigrp

Routing Protocol: EIGRP, instance 1
 Default context AS: 1, Router ID: 192.168.0.22
  Address Family: IPv4
   Default networks not flagged in outgoing updates
   Default networks not accepted from incoming updates
   Distance: internal 90, external 170
   Maximum paths: 4
   EIGRP metric weight K1=1, K2=0, K3=1, K4=0, K5=0
   EIGRP maximum hopcount 100
   EIGRP maximum metric variance 1
   EIGRP NSF: enabled
    NSF-aware route hold timer is 240s
    NSF signal timer is 20s
    NSF converge timer is 120s
    Time since last restart is 01:01:21
   SIA Active timer is 180s
   Interfaces:
    GigabitEthernet0/6/0/0
```

This table describes the significant fields shown in the display.

Table 58: show protocols Field Descriptions

Field	Descriptions
instance	AS number of the instance. • 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.
AS	AS number of this context. • 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	Address family for which the configuration status is shown.
Default Networks Candidate	Default network acceptance and announcement behavior.
Distance	Administrative distance of EIGRP routes.
Maximum paths	Maximum paths installed in RIB for a route.
Metric Weight	Current metric weights used by EIGRP.
Maximum hopcount	Maximum hop count accepted by EIGRP.

Field	Descriptions
Variance	Metric variance used to find feasible paths for a route.
Route hold time	Time duration for which routes learned from a neighbor are held without deletion while the neighbor is undergoing a graceful restart.
signal time	Nonstop forwarding signal time.
converge time	Nonstop forwarding convergence time.
SIA Active time	Active time period for SIA.
Interfaces	List of interfaces configured for EIGRP.

site-of-origin (EIGRP)

To configure the Site of Origin (SoO) filtering on an Enhanced Interior Gateway Routing Protocol (EIGRP) interface, use the **site-of-origin** command in interface configuration mode. To disable SoO filtering on an interface, use the **no** form of this command.

site-of-origin {*as-number* **:** *number* | *ip-address* **:** *number*}
no site-of-origin

Syntax Description

as-number :	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. The colon is used to separate the autonomous system number and network number.
number	Network number. Range is from 0 to 4294967295 when a 2-byte AS number is used. Range is from 0 to 65535 when a 4-byte AS number is used.
ip-address :	IP address in four-part, dotted-decimal notation. The colon is used to separate the IP address and network number.

Command Default

No default behavior or values

Command Modes

Interface configuration

Command History

Release	Modification
Release 3.7.2	This command was introduced.
Release 3.9.0	Asplain format for 4-byte Autonomous system numbers notation was supported.

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.

An EIGRP process must be capable of retrieving the SoO attribute on routes redistributed from the Border Gateway Protocol (BGP) when required to support complex topologies that include MPLS VPN links between sites with backdoor links.

Use the **site-of-origin** command to set an SoO BGP extended community attribute that is used to identify routes that have originated from a site so that the readvertisement of that prefix back to the source site can be prevented. The SoO extended community uniquely identifies the site from which a provider edge (PE) router has learned a route.

Task ID

Task ID	Operations
eigrp	read, write

Examples

The following example shows how to configure SoO filtering on an EIGRP interface:

```
RP/0/RSP0/CPU0:router(config)# router eigrp 1
RP/0/RSP0/CPU0:router(config-eigrp) vrf customer1
RP/0/RSP0/CPU0:router(config-eigrp-vrf) address-family ipv4
RP/0/RSP0/CPU0:router(config-eigrp-vrf-af)# interface GigabitEthernet 0/1/0/0
RP/0/RSP0/CPU0:router(config-eigrp--vrf-af-if)# site-of-origin 10.0.0.1:20
```

split-horizon disable (EIGRP)

To disable split horizon for an Enhanced Interior Gateway Routing Protocol (EIGRP) process, use the **split-horizon disable** command in interface configuration mode. To enable split horizon, use the **no** form of this command.

split-horizon disable
no split-horizon disable

Syntax Description

This command has no keywords or arguments.

Command Default

Split horizon is enabled for an EIGRP process.

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 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
eigrp	read, write

Examples

The following example shows how to disable split horizon an a GigabitEthernet link:

```
RP/0/RSP0/CPU0:router(config)# router eigrp 1
RP/0/RSP0/CPU0:router(config-eigrp) address-family ipv4
RP/0/RSP0/CPU0:router(config-eigrp-af)# interface GigabitEthernet 0/1/0/0
RP/0/RSP0/CPU0:router(config-eigrp-af-if)# split-horizon disable
```

stub (EIGRP)

To configure a router as a stub for Enhanced Interior Gateway Routing Protocol (EIGRP), use the **stub** command in the appropriate configuration mode. To disable this function, use the **no** form of this command.

stub [**receive-only** | [**connected**] [**redistributed**] [**static**] [**summary**]]
no stub [**receive-only** | [**connected**] [**redistributed**] [**static**] [**summary**]]

Syntax Description

receive-only	(Optional) Sets the router as a receive-only neighbor.
connected	(Optional) Advertises connected routes.
redistributed	(Optional) Advertises redistributed routes from other protocols and autonomous systems.
static	(Optional) Advertises static routes.
summary	(Optional) Advertises summary routes.

Command Default

Stub routing is disabled.

When stub routing is specified, connected and summary routes are advertised by default.

Command Modes

IPv4 address family configuration

IPv4 VRF address family 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **stub** command to configure a router as a stub in which the router directs all IP traffic to a distribution router.

The **stub** command can be modified with several options, and these options can be used in any combination except for the **receive-only** keyword.

The **receive-only** keyword restricts the router from sharing any of its routes with any other router in that EIGRP autonomous system and does not permit any other option to be specified because it prevents any type of route from being sent. The four other optional keywords (**connected** , **static** , **summary** , and **redistributed**) can be used in any combination but cannot be used with the **receive-only** keyword. If any of these four keywords is used with the **stub** command, only the route types specified by the particular keyword or keywords are sent. Route types specified by the nonused keyword or keywords are not sent.

The **connected** keyword permits EIGRP stub routing to send connected routes. If all the connected routes are not covered by EIGRP interfaces, it may be necessary to redistribute connected routes with the **redistribute connected** command under the EIGRP process. This option is enabled by default.

The **static** keyword permits EIGRP stub routing to send static routes. Without the configuration of this option, EIGRP does not send any static routes. You may still need to redistribute static routes with the **redistribute static** command.

The **summary** keyword permits EIGRP stub routing to send summary routes. Summary routes can be created manually with the **summary address** command or automatically at a major network border router with the **auto-summary** command enabled. This option is enabled by default.

The **redistributed** keyword permits EIGRP stub routing to send other routing protocols and autonomous systems. Without the configuration of this option, EIGRP does not advertise redistributed routes.

Task ID

Task ID	Operations
eigrp	read, write

Examples

The following example shows how to configure, as a stub, the router that advertises connected and summary routes:

```
RP/0/RSP0/CPU0:router(config)# router eigrp 1
RP/0/RSP0/CPU0:router(config-eigrp) address-family ipv4
RP/0/RSP0/CPU0:router(config-eigrp-af)# stub
```

The following example shows how to configure the router as a receive-only neighbor (connected, summary, and static routes are not sent):

```
RP/0/RSP0/CPU0:router(config)# router eigrp 1
RP/0/RSP0/CPU0:router(config-eigrp) address-family ipv4
RP/0/RSP0/CPU0:router(config-eigrp-af)# stub receive-only
```

Related Commands

Command	Description
redistribute (EIGRP), on page 760	Redistributes routes from one routing domain into EIGRP.
summary-address (EIGRP), on page 791	Configures a summary aggregate address for the specified EIGRP interface.
auto-summary (EIGRP), on page 717	Allows automatic summarization of subnet routes into network-level routes for an EIGRP process.

summary-address (EIGRP)

To configure a summary aggregate address for the specified Enhanced Interior Gateway Routing Protocol (EIGRP) interface, use the **summary-address** command in interface configuration mode. To disable a configuration, use the **no** form of this command.

summary-address *ip-address* {*/lengthmask*} [*admin-distance*]
no summary-address *ip-address* {*/lengthmask*}

Syntax Description

ip-address	The IP address argument specifies the summary IP address to apply to an interface in four-part, dotted-decimal notation.
/ length	Prefix length, which can be indicated as a slash (/) and number. For example, /8 indicates that the first eight bits in the IP prefix are network bits. If *length* is used, the slash is required.
mask	IP address mask.
admin-distance	(Optional) Administrative distance. A value from 1 to 255.

Command Default

An administrative distance of 5 is applied to EIGRP summary routes.

No summary addresses are predefined.

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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

The **summary-address** command is used to configure interface-level address summarization. EIGRP summary routes are given an administrative distance of 5. The administrative distance is used to advertise a summary without installing it in the routing table.

By default, EIGRP summarizes subnet routes to the network level. The **no auto-summary** command can be entered to configure subnet level summarization.

Task ID

Task ID	Operations
eigrp	read, write

Examples

The following example shows how to configure an administrative distance of 95 on an EIGRP interface for the 192.168.0.0/16 summary address:

```
RP/0/RSP0/CPU0:router(config)# router eigrp 1
RP/0/RSP0/CPU0:router(config-eigrp) address-family ipv4
RP/0/RSP0/CPU0:router(config-eigrp-af)# interface GigabitEthernet 0/1/0/0
RP/0/RSP0/CPU0:router(config-eigrp-af-if)# summary-address 192.168.0.0/16 95
```

Related Commands

Command	Description
auto-summary (EIGRP), on page 717	Allows automatic summarization of subnet routes into network-level routes for an EIGRP process.

timers active-time

To adjust the Enhanced Interior Gateway Routing Protocol (EIGRP) routing wait time, use the **timers active-time** command in the appropriate configuration mode. To disable this function, use the **timers active-time no** form of the command.

timers active-time [*time-limit* | **disabled**]
no timers active-time

Syntax Description

time-limit	Active time limit (in minutes). Range is from 1 to 4294967295 minutes.
disabled	Disables the timers and permits the routing wait time to remain active indefinitely.

Command Default

Disabled

Command Modes

IPv4 address family configuration

IPv4 VRF address family 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **timers active-time** command to control the time the router waits (after query is sent) before declaring the route to be in the stuck in active (SIA) state.

Task ID

Task ID	Operations
eigrp	read, write

Examples

The following example shows how to configure an indefinite routing wait time on the specified EIGRP route:

```
RP/0/RSP0/CPU0:router(config)# router eigrp 1
RP/0/RSP0/CPU0:router(config-eigrp) address-family ipv4
RP/0/RSP0/CPU0:router(config-eigrp-af)# timers active-time disabled
```

timers nsf route-hold (EIGRP)

To set the timer that determines how long an NSF-aware Enhanced Interior Gateway Routing Protocol (EIGRP) router holds routes for an inactive peer, use the **timers nsf route-hold** command in the appropriate configuration mode. To return the route hold timer to the default value, use the **no** form of this command.

timers nsf route-hold *seconds*
no timers nsf route-hold

Syntax Description

seconds	Time, in seconds, that EIGRP holds routes for an inactive peer. Range is from 20 to 6000 seconds.

Command Default

EIGRP NSF awareness is enabled.

seconds :240

Command Modes

IPv4 address family configuration

IPv4 VRF address family configuration

Command History

Release	Modification
Release 3.7.2	This command was introduced.
Release 4.1.1	The default route hold time was changed from 240 seconds to 480 seconds.

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.

Use the **timers nsf route-hold** command to set the maximum period of time that the NSF-aware router holds known routes for an NSF-capable neighbor during a switchover operation or a well-known failure condition. The route hold timer is configurable so that you can tune network performance and avoid undesired effects, such as packet loss on routes if the switchover operation takes too much time. When this timer expires, the NSF-aware router scans the topology table and discards any stale routes, allowing EIGRP peers to find alternate routes instead of waiting during a long switchover operation.

Task ID

Task ID	Operations
eigrp	read, write

Examples

The following example shows how to set the route hold timer value for an NSF-aware router to 2 minutes (120 seconds):

```
RP/0/RSP0/CPU0:router(config)# router eigrp 1
RP/0/RSP0/CPU0:router(config-eigrp) address-family ipv4
RP/0/RSP0/CPU0:router(config-eigrp-af)# timers nsf route-hold 120
```

variance

To control load balancing in an Enhanced Interior Gateway Routing Protocol (EIGRP)-based internetwork, use the **variance** command in the appropriate configuration mode. To reset the variance to the default value, use the **no** form of this command.

variance *multiplier*
no variance

Syntax Description

multiplier	Metric value used for load balancing. Range is from 1 to 128.

Command Default

multiplier : 1 (equal-cost load balancing)

Command Modes

IPv4 address family configuration

IPv4 VRF address family 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **variance** command to set a variance on the EIGRP router so that the router can determine the feasibility of a potential route. A route is feasible if the next router in the path is closer to the destination than the current router and the metric for the entire path is within the variance. Only paths that are feasible can be used for load balancing and included in the routing table.

If the following two conditions are met, the route is considered feasible and can be added to the routing table:

1. The local best metric must be greater than the metric learned from the next router.
2. The multiplier times the local best metric for the destination must be greater than or equal to the metric through the next router.

Task ID

Task ID	Operations
eigrp	read, write

Examples

The following example shows how to set a variance of 4:

```
RP/0/RSP0/CPU0:router(config)# router eigrp 1
RP/0/RSP0/CPU0:router(config-eigrp) address-family ipv4
RP/0/RSP0/CPU0:router(config-eigrp-af)# variance 4
```

vrf (EIGRP)

To define a VPN routing and forwarding (VRF) instance and enter VRF configuration mode, use the **vrf** command in router configuration mode. To remove a VRF instance, use the **no** form of this command.

vrf *vrf-name*
no vrf *vrf-name*

Syntax Description

vrf-name	VPN routing and forwarding instance.

Command Default

No VRFs are defined.

Command Modes

Router 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

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.

From VRF configuration mode, you must enter address family configuration mode and then issue commands, such as the **auto-summary** command.

A single EIGRP routing process can support multiple VRFs. The number of VRFs that can be configured is limited by available system resources on the router, which is determined by the number of VRFs, running processes, and available memory. However, only a single VRF can be supported by each VPN. Redistribution between different VRFs is not supported.

MPLS VPN support between PE and customer edge (CE) routers is configured only on PE routers that provide VPN services over the service provider backbone. The customer site does not require any changes to equipment or configurations to support the EIGRP VPN. Typically, a metric must be configured for routes to be advertised to the CE router. The metric can be configured under the route-policy for the **redistribute** protocol command or configured with the **default-metric** command.

You must remove IPv4/IPv6 addresses from an interface prior to assigning, removing, or changing a VRF on an IP interface. If this is not done in advance, any attempt to change the VRF on an IP interface is rejected.

Task ID

Task ID	Operations
eigrp	read, write

Examples

The following example shows how to enter IPv4 VRF address family configuration mode and identify EIGRP commands that can be issued from that mode.

```
RP/0/RSP0/CPU0:router(config)# router eigrp 1
RP/0/RSP0/CPU0:router(config-eigrp)# vrf vpn-1
RP/0/RSP0/CPU0:router(config-eigrp-vrf)# address-family ipv4
RP/0/RSP0/CPU0:router(config-eigrp-vrf-af)# ?

  auto-summary            Auto summarisation
  autonomous-system       Set the autonomous system of VRF
  commit                  Commit the configuration changes to running
  default-information     Handling of default route
  default-metric          Set metric of redistributed routes
  describe                Describe a command without taking real actions
  distance                Set distance for EIGRP routes
  do                      Run an exec command
  exit                    Exit from this submode
  interface               EIGRP interface configuration submode
  log-neighbor-changes    Enable/Disable EIGRP neighbor logging
  log-neighbor-warnings   Enable/Disable EIGRP neighbor warnings
  maximum-paths           Maximum paths
  maximum-prefix          Maximum number of IP prefixes acceptable in aggregate
  metric                  Modify EIGRP routing metrics and parameters
  neighbor                Neighbor prefix limits configuration
  no                      Negate a command or set its defaults
  redistribute            Redistribute another protocol
  route-policy            Configure inbound/outbound policies
  router-id               Set router ID
  show                    Show contents of configuration
  stub                    EIGRP stub
  timers                  Configure EIGRP timers
  variance                Control load balancing variance
```

Related Commands

Command	Description
default-metric (EIGRP), on page 730	Sets metrics for an EIGRP.
redistribute (EIGRP), on page 760	Injects routes from one routing domain EIGRP.



IS-IS Commands

This module describes the commands used to configure and monitor the Intermediate System-to-Intermediate System (IS-IS) protocol on Cisco ASR 9000 Series Aggregation Services Routers .

For detailed information about IS-IS concepts, configuration tasks, and examples, see the *Implementing IS-IS on* Cisco ASR 9000 Series Router module in the *Routing Configuration Guide for Cisco ASR 9000 Series Routers*.

- address-family (IS-IS), on page 803
- address-family multicast topology (IS-IS), on page 805
- adjacency-check disable, on page 807
- advertise passive-only, on page 809
- adjacency stagger (IS-IS), on page 810
- algorithm 0, on page 812
- attached-bit receive ignore, on page 813
- attached-bit send, on page 814
- authentication-check disable, on page 816
- circuit-type, on page 817
- clear isis process, on page 819
- clear isis route, on page 820
- clear isis statistics, on page 822
- cost-fallback anomaly delay, on page 824
- csnp-interval, on page 826
- default-information originate (IS-IS), on page 828
- disable (IS-IS), on page 830
- distance (IS-IS), on page 831
- distribute-list in (IS-IS), on page 833
- fast-reroute per-link (IS-IS), on page 835
- fast-reroute per-prefix (IS-IS), on page 837
- fast-reroute per-link priority-limit (IS-IS), on page 839
- fast-reroute per-link use-candidate-only (IS-IS), on page 841
- fast-reroute per-prefix load-sharing disable (IS-IS), on page 842
- fast-reroute per-prefix tiebreaker (IS-IS), on page 843
- fast-reroute per-prefix use-candidate-only (IS-IS), on page 845
- hello-interval (IS-IS), on page 846
- hello-multiplier, on page 848

- hello-padding, on page 850
- hello-password, on page 852
- hello-password keychain, on page 854
- hello-password accept, on page 856
- hostname dynamic disable, on page 858
- ignore-lsp-errors, on page 859
- instance-id, on page 860
- interface (IS-IS), on page 861
- ipfrr lfa, on page 863
- ipfrr lfa exclude interface, on page 865
- ispf, on page 867
- is-type, on page 868
- link-down fast-detect, on page 870
- link-group, on page 871
- log adjacency changes (IS-IS), on page 872
- log pdu drops, on page 873
- lsp fast-flood threshold, on page 874
- lsp-fast-flooding, on page 875
- lsp-gen-interval, on page 877
- lsp-interval, on page 879
- lsp-mtu, on page 880
- lsp-password, on page 882
- lsp-password accept, on page 884
- lsp-refresh-interval, on page 886
- maximum-paths (IS-IS), on page 888
- maximum-redistributed-prefixes (IS-IS), on page 889
- maximum-propagated-prefixes (IS-IS), on page 890
- max-lsp-lifetime, on page 891
- max-lsp-tx, on page 892
- max-metric, on page 893
- metric fallback anomaly, on page 895
- mesh-group (IS-IS), on page 897
- metric (IS-IS), on page 899
- metric-style narrow, on page 902
- metric-style transition, on page 904
- metric-style wide, on page 906
- microloop avoidance (IS-IS), on page 908
- min-lsp-arrivaltime, on page 910
- mpls ldp auto-config , on page 912
- mpls ldp sync (IS-IS), on page 913
- mpls traffic-eng (IS-IS), on page 915
- mpls traffic-eng multicast-intact (IS-IS), on page 917
- mpls traffic-eng path-selection ignore overload, on page 918
- mpls traffic-eng router-id (IS-IS), on page 920
- net, on page 922
- nsf (IS-IS), on page 924

- nsf interface-expires, on page 926
- nsf interface-timer, on page 928
- nsf lifetime (IS-IS), on page 930
- oor-set-overload-bit disable, on page 931
- passive (IS-IS), on page 932
- point-to-point, on page 933
- priority (IS-IS), on page 934
- propagate level, on page 936
- protocol shutdown (IS-IS), on page 938
- psnp-interval, on page 939
- redistribute (IS-IS), on page 940
- remote-psnp-delay, on page 944
- retransmit-interval (IS-IS), on page 946
- retransmit-throttle-interval, on page 948
- router isis, on page 950
- router isis static neighbor, on page 951
- route source first-hop, on page 952
- set default-isis-instance, on page 953
- set-overload-bit, on page 955
- set-attached-bit, on page 957
- show isis, on page 959
- show isis adjacency, on page 961
- show isis adjacency-log, on page 963
- show isis checkpoint adjacency, on page 965
- show isis checkpoint interface, on page 967
- show isis checkpoint lsp, on page 969
- show isis database, on page 971
- show isis database-log, on page 972
- show isis fast-reroute, on page 974
- show isis hostname, on page 977
- show isis interface, on page 979
- show isis lsp-log, on page 983
- show isis mesh-group, on page 985
- show isis mpls traffic-eng adjacency-log, on page 987
- show isis mpls traffic-eng advertisements, on page 989
- show isis neighbors, on page 992
- show isis protocol, on page 995
- show isis route, on page 998
- show isis spf-log, on page 1001
- show isis statistics, on page 1008
- show isis topology, on page 1012
- show protocols (IS-IS), on page 1015
- shutdown (IS-IS), on page 1019
- single-topology, on page 1020
- snmp-server traps isis, on page 1021
- spf-interval, on page 1022

- spf-interval ietf, on page 1024
- spf prefix-priority (IS-IS), on page 1026
- summary-prefix (IS-IS), on page 1028
- suppressed, on page 1030
- tag (IS-IS), on page 1031
- topology-id, on page 1032
- trace (IS-IS), on page 1033
- ucmp (IS-IS), on page 1034
- ucmp delay-interval, on page 1036
- ucmp exclude interface, on page 1038

address-family (IS-IS)

To enter address family configuration mode for configuring Intermediate System-to-Intermediate System (IS-IS) routing that use standard IP Version 4 (IPv4) and IP Version 6 (IPv6) address prefixes, use the **address-family** command in router configuration or interface configuration mode. To disable support for an address family, use the **no** form of this command.

address-family {**ipv4** | **ipv6**} {**unicast** | **multicast**}
no address-family {**ipv4** | **ipv6**} {**unicast** | **multicast**}

Syntax Description

ipv4	Specifies IPv4 address prefixes.
ipv6	Specifies IPv6 address prefixes.
unicast	Specifies unicast address prefixes.
multicast	Specifies multicast address prefixes.

Command Default

An address family is not specified. The default subaddress family (SAFI) is unicast.

Command Modes

Router configuration

Interface configuration

Command History

Release	Modification
Release 3.7.2	This command was introduced.
Release 3.9.0	Support for IPv6 was added.

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.

Use the **address family** command to place the router or interface in address family configuration mode. In router address family configuration mode, you can configure routing that uses standard IPv4 or IPv6 address prefixes. An address family must be specified in interface configuration mode. In interface address family configuration mode, you can alter interface parameters for IPv4or IPv6.

You must specify an address family in order to configure parameters that pertain to a single address family.

Task ID

Task ID	Operations
isis	read, write

Examples

The following example shows how to configure the IS-IS router process with IPv4 unicast address prefixes:

```
RP/0/RSP0/CPU0:router(config)# router isis isp
RP/0/RSP0/CPU0:router(config-isis)# interface gigabitEthernet 0/1/0/0
RP/0/RSP0/CPU0:router(config-isis-if)# address-family ipv4 unicast
RP/0/RSP0/CPU0:router(config-isis-if-af)#
```

address-family multicast topology (IS-IS)

To enable a multicast topology when configuring Intermediate System-to-Intermediate System (IS-IS) routing (or to place a given topology within the IS-IS interface), use the **address-family multicast topology** command with either IPv4 or IPv6 address prefix in the appropriate configuration mode. To disable a multicast topology in IS-IS, use the **no** form of this command.

address-family {**ipv4** | **ipv6**} **multicast topology** *topo-name* [**maximum prefix** *prefix-limit*]
no address-family

Syntax Description

ipv4	Specifies IPv4 address prefixes.
ipv6	Specifies IPv6 address prefixes.
multicast	Specifies multicast address prefixes.
topology *topo-name*	Specifies the name of the topology.
maximum prefix	Specifies maximum number of prefixes that a routing table can have.
prefix-limit	Maximum number of prefixes. Range is from 32 to 2,000,000.

Command Default

An address family for multicast topology is not specified. The default subaddress family (SAFI) is unicast.

Command Modes

Router configuration

Interface configuration

Command History

Release	Modification
Release 3.7.2	This command was introduced.
Release 3.9.0	Support for IPv6 was added.

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.

Use the **address family multicast topology** command to place the router or interface in address family configuration mode. In router address family configuration mode, you can associate an IS-IS topology ID with the topology you have created to add connected and local routes to a specific routing table.

Task ID

Task ID	Operations
isis	read, write

Examples

The following example shows how to configure the IS-IS router topology with an IPv4 multicast address prefix:

```
RP/0/RSP0/CPU0:router(config)# router isis isp
RP/0/RSP0/CPU0:router(config-isis)# address-family ipv6 multicast topology green
RP/0/RSP0/CPU0:router(config-isis-af)#
```

or

```
RP/0/RSP0/CPU0:router(config)# router isis isp
RP/0/RSP0/CPU0:router(config-isis)#  interface gigabitethernet 0/3/0/0
RP/0/RSP0/CPU0:router(config-isis-if)# address-family ipv4 multicast topology green
RP/0/RSP0/CPU0:router(config-isis-if-af)#
```

Related Commands

Command	Description
topology-id, on page 1032	Associates a topology ID with a named IS-IS topology to differentiate topologies in the domain.

adjacency-check disable

To suppress Intermediate System-to-Intermediate System (IS-IS) IP Version 4 (IPv4) or IP Version 6 (IPv6) protocol-support consistency checks that are performed prior to forming adjacencies on hello packets, use the **adjacency-check disable** command in address family configuration mode. To remove this function, use the **no** form of this command.

adjacency-check disable
no adjacency-check disable

Command Default

Adjacency check is enabled

Command Modes

Address family configuration

Command History

Release	Modification
Release 3.7.2	This command was introduced.
Release 3.9.0	Support was added for ipv6.

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.

IS-IS performs consistency checks on hello packets and forms an adjacency only with a neighboring router that supports the same set of protocols. A router running IS-IS for both IPv4 and IPv6 does not form an adjacency with a router running IS-IS for IPv4 only.

Use the **adjacency-check disable** command to suppress the consistency checks for IPv6 IS-IS and allow an IPv4 IS-IS router to form an adjacency with a router running IPv4 IS-IS and IPv6. IS-IS never forms an adjacency between a router running IPv4 IS-IS only and a router running IPv6 only.

In addition, the **adjacency-check disable** command suppresses the IPv4 or IPv6 subnet consistency check and allows IS-IS to form an adjacency with other routers regardless of whether they have an IPv4 or IPv6 subnet in common.

Task ID

Task ID	Operations
isis	read, write

Examples

The command in the following example disables the adjacency checks:

The following example shows how the network administrator introduces IPv6 into an existing IPv4 IS-IS network and ensures that the checking of hello packet checks from adjacent neighbors is disabled until all neighbor routers are configured to use IPv6:

```
RP/0/RSP0/CPU0:router(config)# router isis isp
RP/0/RSP0/CPU0:router(config-isis)# address-family ipv6 |ipv4
```

```
RP/0/RSP0/CPU0:router(config-isis-af)# adjacency-check disable
```

advertise passive-only

To configure IS-IS to advertise only prefixes that belong to passive interfaces, use the **advertise-passive-only** command in ISIS address family configuration mode. To disable advertisement only prefixes that belong to passive interfaces, use the **no** form of this command.

advertise passive-only
no advertise passive-only

Syntax Description

This command has no keywords or arguments.

Command Default

None

Command Modes

IPv4 unicast address family configuration

IPv4 multicast address family configuration

IPv6 unicast address family configuration

IPv6 multicast address family configuration

Command History

Release	Modification
Release 4.2.0	This command was introduced.
Release 5.0.0	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
isis	read, write

This example shows how to configure IS-IS to advertise only prefixes that belong to passive interfaces.

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#router isis isp
RP/0/RSP0/CPU0:router(config-isis)#address-family ipv4 unicast
RP/0/RSP0/CPU0:router(config-isis-af)#advertise passive-only
```

adjacency stagger (IS-IS)

To configure staggering of IS-IS adjacency during reload, process restart, and process clear, use the **adjacency stagger** command in router configuration mode. To turn off adjacency staggering, either use the **disable** keyword or use the **no** form of this command.

adjacency stagger {**disable** | *initial-num-nbr* *max-num-nbr*}
no adjacency stagger

disable	Disables adjacency staggering.
initial-num-nbr	The initial number of simultaneous neighbors allowed to form adjacency to FULL in any area to bring up to FULL after a router reload, IS-IS process restart, or IS-IS process clear. Range is 1-65535. Default is 2.
max-num-nbr	The subsequent number of simultaneous neighbors allowed to form adjacency, per IS-IS instance, after the initial set of IS-IS neighbors have become FULL. Range is 1-65535. Default is 64.

Command Default

IS-IS adjacency staggering is enabled.

Command Modes

Router configuration

Table 59: Command History

Release	Modification
Release 6.3.1	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.

Staggering of the IS-IS adjacency during reload, process restart (without NSR or graceful-restart), and process clear reduces the overall adjacency convergence time.

Initially, allow 2 (configurable) neighbors to form adjacency to FULL per area. After the first adjacency reaches FULL, up to 64 (configurable) neighbors can form adjacency simultaneously for the IS-IS instance (all areas). However, areas without any FULL adjacency is restricted by the initial area limit.



Note Adjacency stagger and IS-IS nonstop forwarding (NSF) are mutually exclusive. Adjacency stagger is not activated if nonstop forwarding (NSF) is configured in the router along with IS-IS configuration.

Table 60: Task ID

Task ID	Operations
IS-IS	read, write

The following example shows how to configure adjacency stagger for a 2 neighbors initially and for a maximum of 32 neighbors:

```
Router# configure
Router(config)# router isis 1
Router(config-isis)# adjacency stagger 2 32
```

algorithm 0

The **algorithm 0** command allows you to override the maximum-paths configuration for an address-family and IS-IS instance pair. The **maximum-paths** subcommand under **algorithm 0** configuration block will be applied to the standard SPF algorithm of that IS-IS instance.

algorithm 0
address-family { **ipv4** | **ipv6** } **unicast**
maximum-paths *maximum-paths*

Syntax Description

address-family {ipv4 | ipv6}	Configures the address family for IS-IS.
unicast	Specifies unicast address prefixes.
maximum-paths	Specifies the maximum number of parallel routes that an IP routing protocol installs in the routing table. The range is 1 to 64.

Command Modes

IS-IS interface

Command History

Release	Modification
Release 7.8.1	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

Task ID

Task ID	Operations
isis	read, write

attached-bit receive ignore

To ignore the attached bit in a received Level 1 link-state packet (LSP), use the **attached-bit receive ignore** command in address family configuration mode. To remove the **attached-bit receive ignore** command from the configuration file and restore the system to its default condition, use the **no** form of this command.

attached-bit receive ignore
no attached-bit receive ignore

Command Default

The attached bit is set in the LSP.

Command Modes

Address family configuration

Command History

Release	Modification
Release 3.8.0	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
isis	read, write

Examples

The following example shows how to configure to ignore the attached bit in a received LSP:

```
RP/0/RSP0/CPU0:router(config)# router isis isp1
RP/0/RSP0/CPU0:router(config-isis)# address-family ipv4 unicast
RP/0/RSP0/CPU0:router(config-isis-af)# attached-bit receive ignore
```

Related Commands

Command	Description
attached-bit send, on page 814	Configures an Intermediate System-to-Intermediate System (IS-IS) instance with an attached bit in the Level 1 link-state packet (LSP).

attached-bit send

To configure an Intermediate System-to-Intermediate System (IS-IS) instance with an attached bit in the Level 1 link-state packet (LSP), use the **attached-bit send** command in address family configuration mode. To remove the **attached-bit send** command from the configuration file and restore the system to its default condition, use the **no** form of this command.

attached-bit send {**always-set** | **never-set**}
no attached-bit send {**always-set** | **never-set**}

Syntax Description

always-set	Specifies to always set the attached bit in the LSP.
never-set	Specifies to never set the attached bit in the LSP.

Command Default

The attached bit is not forced to be set or unset in the LSP.

Command Modes

Address family configuration

Command History

Release	Modification
Release 3.8.0	This command was introduced and replaces the set-attached-bit, on page 957 command.

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.

Use the **attached-bit send** command to set an IS-IS instance with an attached bit in the Level 1 LSP that allows another IS-IS instance to redistribute Level 2 topology. The attached bit is used when the Level 2 connectivity from another IS-IS instance is advertised by the Level 1 attached bit.

Cisco IOS XR software does not support multiple Level 1 areas in a single IS-IS routing instance; however the equivalent functionality is achieved by redistribution of routes between two IS-IS instances by using the redistribute (IS-IS), on page 940 command.

The attached bit is configured for a specific address family only if the **single-topology** command is not configured.

Note If connectivity for the Level 2 instance is lost, the attached bit in the Level 1 instance LSP continues sending traffic to the Level 2 instance and causes the traffic to be dropped.

Task ID

Task ID	Operations
isis	read, write

Examples

The following example shows how to configure an Intermediate System-to-Intermediate System (IS-IS) instance with an attached bit:

```
RP/0/RSP0/CPU0:router(config)# router isis isp1
RP/0/RSP0/CPU0:router(config-isis)# address-family ipv4 unicast
RP/0/RSP0/CPU0:router(config-isis-af)# attached-bit send always-set
```

Related Commands

Command	Description
redistribute (IS-IS), on page 940	Redistribute routes from one routing protocol into Intermediate System-to-Intermediate System (IS-IS).
single-topology, on page 1020	Configures the link topology for IPv4 when IPv6 is configured.

authentication-check disable

To suppress Intermediate System-to-Intermediate System (IS-IS) authentication check, use the **authentication-check disable** command in configuration mode. To remove this function, use the **no** form of this command.

authentication-check disable
no authentication-check disable

Command Default

Authentication check is enabled

Command Modes

Global Configuration mode

Command History

Release	Modification
Release 7.7.1	This command was introduced.

Task ID

Task ID	Operations
isis	read, write

Examples

The following example shows how to disable authentication check for IS-IS.

```
Router# configure
Router(config)# router isis 1
Router(config)# authentication-check disable
Router(config)# commit
```

circuit-type

To configure the type of adjacency used for the Intermediate System-to-Intermediate System (IS-IS) protocol, use the **circuit-type** command in interface configuration mode. To reset the circuit type to Level l and Level 2, use the **no** form of this command.

circuit-type {**level-1** | **level-1-2** | **level-2-only**}
no circuit-type

Syntax Description

level-1	Establishes only Level 1 adjacencies over an interface.
level-1-2	Establishes both Level 1 and Level 2 adjacencies, if possible.
level-2-only	Establishes only Level 2 adjacencies over an interface.

Command Default

Default adjacency types are Level 1 and Level 2 adjacencies.

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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Adjacencies may not be established even if allowed by the **circuit-type** command. The proper way to establish adjacencies is to configure a router as a Level 1, Level 1 and Level 2, or Level 2-only system using the is-type, on page 868 command. Only on networking devices that are between areas (Level 1 and Level 2 networking devices) should you configure some interfaces to be Level 2-only to prevent wasting bandwidth by sending out unused Level 1 hello packets. Remember that on point-to-point interfaces, the Level 1 and Level 2 hello packets are in the same packet.

Task ID

Task ID	Operations
isis	read, write

Examples

The following example shows how to configure a Level 1 adjacency with its neighbor on GigabitEthernetinterface 0/2/0/0 and Level 2 adjacencies with all Level 2-capable routers on GigabitEthernet interface 0/5/0/2:

```
RP/0/RSP0/CPU0:router(config)# router isis isp
RP/0/RSP0/CPU0:router(config-isis)# is-type level-1-2
RP/0/RSP0/CPU0:router(config-isis)# interface GigabitEthernet 0/2/0/0
RP/0/RSP0/CPU0:router(config-isis-if)# circuit-type level-1
```

```
RP/0/RSP0/CPU0:router(config-isis-if)# exit
RP/0/RSP0/CPU0:router(config-isis)# interface GigabitEthernet 0/5/0/2
RP/0/RSP0/CPU0:router(config-isis-if)# circuit-type level-2-only
```

In this example, only Level 2 adjacencies are established because the **is-type** command is configured:

```
RP/0/RSP0/CPU0:router(config)# router isis isp
RP/0/RSP0/CPU0:router(config-isis)# is-type level-2-only
RP/0/RSP0/CPU0:router(config-isis)# interface GigabitEthernet 0/2/0/0
RP/0/RSP0/CPU0:router(config-isis-if)# circuit-type level-1-2
```

Related Commands

Command	Description
is-type, on page 868	Configures the routing level for an instance of the IS-IS routing process.
net, on page 922	Configures an IS-IS NET for the routing process.

clear isis process

To clear the link-state packet (LSP) database and adjacency database sessions for an Intermediate System-to-Intermediate System (IS-IS) instance or all IS-IS instances, use the **clear isis process** command in EXEC configuration mode.

clear isis [**instance** *instance-id*] **process**

Syntax Description

instance *instance-id*	(Optional) Specifies IS-IS sessions for the specified IS-IS instance only. • The *instance-id* argument is the instance identifier (alphanumeric) defined by the **router isis** command.

Command Default

No default behavior or values

Command Modes

EXEC 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **clear isis process** command without any keyword to clear all the IS-IS instances. Add the **instance** *instance-id* keyword and argument to clear the specified IS-IS instance.

Task ID

Task ID	Operations
isis	read, write

Examples

The following example shows the IS-IS LSP database and adjacency sessions being cleared for instance 1:

```
RP/0/RSP0/CPU0:router# clear isis instance 1 process
```

Related Commands

Command	Description
show isis database, on page 971	Displays the IS-IS link-state database.
show isis neighbors, on page 992	Displays information about IS-IS neighbors.

clear isis route

To clear the Intermediate System-to-Intermediate System (IS-IS) routes in a topology, use the **clear isis route** command in EXEC configuration mode.

clear isis [**instance** *instance-id*] {**afi-all** | **ipv4** | **ipv6**} {**unicast** | **multicast** | **safi-all**} [**topology** *topo-name*] **route**

Syntax Description

instance *instance-id*	(Optional) Specifies IS-IS sessions for the specified IS-IS instance only. • The *instance-id* argument is the instance identifier (alphanumeric) defined by the **router isis** command.
afi-all	Specifies IP Version 4 (IPv4) and IP Version 6 (IPv6) address prefixes.
ipv4	Specifies IPv4 address prefixes.
ipv6	Specifies IPv6 address prefixes.
unicast	Specifies unicast address prefixes.
multicast	Specifies multicast address prefixes.
safi-all	Specifies all secondary address prefixes.
topology *topo-name*	(Optional) Specifies topology table information and name of the topology table.

Command Default

No default behavior or value

Command Modes

EXEC configuration

Command History

Release	Modification
Release 3.7.2	This command was introduced.
Release 3.9.0	Support for IPv6 was added.

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.

Use the **clear isis route** command to clear the routes from the specified topology or all routes in all topologies if no topology is specified.

Task ID

Task ID	Operations
isis	execute
rib	read, write

Task ID	Operations
basic-services	read, write

Examples

The following example shows how to clear the routes with IPv4 unicast address prefixes:

```
RP/0/RSP0/CPU0:router# clear isis ipv4 unicast route
```

Related Commands

Command	Description
show isis database, on page 971	Displays the IS-IS link-state database.
show isis neighbors, on page 992	Displays information about IS-IS neighbors.

clear isis statistics

To clear the Intermediate System-to-Intermediate System (IS-IS) statistics, use the **clear isis statistics** command in EXEC configuration mode.

clear isis [**instance** *instance-id*] **statistics** [*type interface-path-id*]

Syntax Description

instance *instance-id*	(Optional) Clears IS-IS sessions for the specified IS-IS instance only. • The *instance-id* argument is the instance identifier (alphanumeric) defined by the **router isis** command.
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

No default behavior or values

Command Modes

EXEC 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **clear isis statistics** command to clear the information displayed by the **show isis statistics** command.

Task ID

Task ID	Operations
isis	execute
rib	read, write
basic-services	read, write

Examples

The following example shows the IS-IS statistics for a specified interface being cleared:

```
RP/0/RSP0/CPU0:router# clear isis instance 23 statistics
```

Related Commands

Command	Description
show isis statistics, on page 1008	Displays the IS-IS statistics.

cost-fallback anomaly delay

cost-fallback anomaly delay { **igp-metric** | **te-metric** } { **increment** *penalty_value* | **multiplier** *penalty_multiplier* | **value** *penalized_metric* | **disable** }

Syntax Description

increment *penalty_value*	Increment the IGP cost by the specified value. The range is from 1 to 65534.
multiplier*penalty_multiplier*	Multiply the IGP cost by the specified value. The range is from 1 to 255.
value *penalized_metric*	Set the IGP cost to the specified value. The range is from 1 to 65535.

Command Modes

Interface Configuration

Area Configuration

Router Configuration

Command History

Release	Modification
Release 7.4.1	This command was introduced.

Usage Guidelines

Use the **cost-fallback anomaly delay** command to increase the metric when you configure link delay anomaly detection in Segment Routing Performance Measurement (SR-PM).

When you configure link delay anomaly detection, SR-PM sends an anomaly bit (A-bit) when the upper bound of the anomaly threshold is exceeded. When an anomaly is received for an interface, OSPF can increase the IGP metric or TE metric to make the link less attractive for traffic.

You can specify how to increase the metric in the following ways:

- As a factor by which the base metric is multiplied
- As an increment that is added to the base metric
- As a fixed high value (up to 0xFFFFFFFF).

Note Due to protocol limitations, there is a maximum value applied for the metric for all three options:

- IGP metric maximum value: 0xFFFF
- TE metric maximum value: 0xFFFFFFFF

The configuration can be applied at the OSPF interface, area, and instance levels, and is applicable only to non-loopback interfaces. When the configuration is applied at the OSPF interface level, it is applicable to that interface only.

Task ID

Task ID	Operation
ospf	read, write

Example

In the following examples, the base IGP metric values are 10.

This example shows how to increase the IGP and TE metrics for an interface. The IGP metric will be increased by 10 and the TE metric will be multiplied by 10. When an anomalous bit notification comes for interface GigabitEthernet0/2/0/0, the values advertised in the TLVs will be IGP metric 20 and TE metric 100.

```
Router(config)# router ospf 1
Router(config-ospf)# area 0
Router(config-ospf-ar)# interface GigabitEthernet 0/2/0/0
Router(config-ospf-ar-if)# cost-fallback anomaly delay igp-metric increment 10
Router(config-ospf-ar-if)# cost-fallback anomaly delay te-metric multiplier 10
```

This example shows how to set a specific IGP and TE metric value for an interface. When an anomalous bit notification comes for interface GigabitEthernet0/2/0/0, the values advertised in the TLVs will be IGP metric 100 and TE metric 50.

```
Router(config)# router ospf 1
Router(config-ospf)# area 0
Router(config-ospf-ar)# interface GigabitEthernet 0/2/0/0
Router(config-ospf-ar-if)# cost-fallback anomaly delay igp-metric value 100
Router(config-ospf-ar-if)# cost-fallback anomaly delay te-metric value 50
```

csnp-interval

To configure the interval at which periodic complete sequence number PDU (CSNP) packets are sent on broadcast interfaces, use the **csnp-interval** command in interface configuration mode. To restore the default value, use the **no** form of this command.

csnp-interval *seconds* [**level** {**1** | **2**}]
no csnp-interval *seconds* [**level** {**1** | **2**}]

Syntax Description

seconds	Interval (in seconds) of time between transmission of CSNPs on multiaccess networks. This interval applies only for the designated router. Range is 0 to 65535 seconds.
level { **1** | **2** }	(Optional) Specifies the interval of time between transmission of CSNPs for Level 1 or Level 2 independently.

Command Default

seconds : 10 seconds

Both Level 1 and Level 2 are configured if no level is specified.

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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

The **csnp-interval** command applies only to the designated router (DR) for a specified interface. Only DRs send CSNP packets to maintain database synchronization. The CSNP interval can be configured independently for Level 1 and Level 2.

Use of the **csnp-interval** command on point-to-point subinterfaces makes sense only in combination with the IS-IS mesh-group feature.

Task ID

Task ID	Operations
isis	execute
rib	read, write
basic-services	read, write

Examples

The following example shows how to set the CSNP interval for Level 1 to 30 seconds:

```
RP/0/RSP0/CPU0:router(config)# router isis isp
```

```
RP/0/RSP0/CPU0:router(config-isis)# interface GigabitEthernet 0/0/2/0
RP/0/RSP0/CPU0:router(config-isis-if)# csnp-interval 30 level 1
```

default-information originate (IS-IS)

To generate a default route into an Intermediate System-to-Intermediate System (IS-IS) routing domain, use the **default-information originate** command in address family configuration mode. To remove the **default-information originate** command from the configuration file and restore the system to its default condition, use the **no** form of this command.

default-information originate [**route-policy** *route-policy-name*]
no default-information originate [**external** | **route-policy** *route-policy-name*]

Syntax Description

route-policy	(Optional) Defines the conditions for the default route.
route-policy-name	(Optional) Name for the route policy.

Command Default

A default route is not generated into an IS-IS routing domain.

Command Modes

Address family 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

If a router configured with the **default-information originate** command has a route to 0.0.0.0 in the routing table, IS-IS originates an advertisement for 0.0.0.0 in its link-state packets (LSPs).

Without a route policy, the default is advertised only in Level 2 LSPs. For Level 1 routing, there is another process to find the default route, which is to look for the closest Level 1 and Level 2 router. The closest Level 1 and Level 2 router can be found by looking at the attached-bit (ATT) in Level 1 LSPs.

A route policy can be used for two purposes:

- To make the router generate the default route in its Level 1 LSPs.
- To advertise 0.0.0.0/0 conditionally.

Task ID

Task ID	Operations
isis	read, write

Examples

The following example shows how to generate a default external route into an IS-IS domain:

```
RP/0/RSP0/CPU0:router(config)# router isis isp
```

```
RP/0/RSP0/CPU0:router(config-isis)# address-family ipv4 unicast
RP/0/RSP0/CPU0:router(config-isis-af)# default-information originate
```

Related Commands

Command	Description
redistribute (IS-IS), on page 940	Redistributes routes from one routing protocol into Intermediate System-to-Intermediate System (IS-IS).
show isis database, on page 971	Displays the IS-IS link-state database.

disable (IS-IS)

To disable the Intermediate System-to-Intermediate System (IS-IS) topology on a specified interface, use the **disable** command in interface address family configuration mode. To remove this function, use the **no** form of this command.

disable
no disable

Command Default

IS-IS protocol is enabled.

Command Modes

Interface address family 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 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
isis	read, write

Examples

The following example shows how to disable the IS-IS protocol for IPv4 unicast on GigabitEthernet interface 0/1/0/1:

```
RP/0/RSP0/CPU0:router(config)# router isis isp
RP/0/RSP0/CPU0:router(config-isis)# interface GigabitEthernet 0/1/0/1
RP/0/RSP0/CPU0:router(config-isis-if)# address-family ipv4 unicast
RP/0/RSP0/CPU0:router(config-isis-if-af)# disable
```

distance (IS-IS)

To define the administrative distance assigned to routes discovered by the Intermediate System-to-Intermediate System (IS-IS) protocol, use the **distance** command in address family configuration mode. To remove the **distance** command from the configuration file and restore the system to its default condition in which the software removes a distance definition, use the **no** form of this command.

distance *weight* [*prefix maskprefix/length* | [*prefix-list-name*]]
no distance [*weight*] [*prefix maskprefix/length* | [*prefix-list-name*]]

Syntax Description

weight	Administrative distance to be assigned to IS-IS routes. Range is 1 to 255.
prefix	(Optional) The *prefix* argument specifies the IP address in four-part, dotted-decimal notation.
mask	(Optional) IP address mask.
/length	(Optional) The length of the IP 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. Range is 0 to 32 for IPv4 addresses and 0 to 128 for IPv6 addresses.
prefix-list-name	(Optional) List of routes to which administrative distance applies.

Command Default

weight : 115

Command Modes

Address family 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

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 that the routing information source cannot be trusted at all and should be ignored. Weight values are subjective; no quantitative method exists for choosing weight values.

Use the **distance** command to configure the administrative distances applied to IS-IS routes when they are inserted into the Routing Information Base (RIB), and influence the likelihood of these routes being preferred over routes to the same destination addresses discovered by other protocols.

The *address/prefix-length* argument defines to which source router the distance applies. In other words, each IS-IS route is advertised by another router, and that router advertises an address that identifies it. This source address is displayed in the output of the **show isis route detail** command.

The **distance** command applies to the routes advertised by routers whose address matches the specified prefix. The *prefix-list-name* argument can then be used to refine this further so that the **distance** command affects only specific routes.

Task ID

Task ID	Operations
isis	read, write

Examples

In the following example, a distance of 10 is assigned to all routes to 2.0.0.0/8 and 3.0.0.0/8 (or more specific prefixes) that are advertised by routers whose ID is contained in 1.0.0.0/8. A distance of 80 is assigned to all other routes.

```
RP/0/RSP0/CPU0:router(config)# ipv4 prefix-list target_routes
RP/0/RSP0/CPU0:router(config-ipv4_pfx)# permit 2.0.0.0/8
RP/0/RSP0/CPU0:router(config-ipv4_pfx)# permit 3.0.0.0/8
RP/0/RSP0/CPU0:router(config-ipv4_pfx)# deny 0.0.0.0/0
RP/0/RSP0/CPU0:router(config-ipv4_pfx)# exit
RP/0/RSP0/CPU0:router(config)# router isis isp
RP/0/RSP0/CPU0:router(config-isis)# address-family ipv4 unicast
RP/0/RSP0/CPU0:router(config-isis-af)# distance 10 1.0.0.0/8 target_routes
RP/0/RSP0/CPU0:router(config-isis-af)# distance 80
```

Related Commands

Command	Description
router isis, on page 950	Configures the IS-IS routing protocol and specifies an IS-IS instance.
show isis protocol, on page 995	Displays summary information about the IS-IS instance.
show isis route, on page 998 **detail**	Displays link-state packet (LSP) details.

distribute-list in (IS-IS)

To filter the routes that Intermediate System-to-Intermediate System (IS-IS) installs in the Routing Information Base (RIB), use the **distribute-list in** command in the appropriate configuration mode. To change or remove the filter, use the **no** form of this command.

distribute-list { {**prefix-list** *prefix-list-name* | **route-policy** *route-policy-name*} } **in**
no distribute-list { {**prefix-list** *prefix-list-name* | **route-policy** *route-policy-name*} } **in**

Syntax Description

prefix-list	IP Version 4 (IPv4) or IP Version 6 (IPv6) prefix list. The list defines which prefixes are to be installed and which are to be omitted when installing routes in the RIB.
route-policy	Specifies the route-policy to filter IS-IS routes.
in	Applies the prefix list or route-policy to incoming routing updates.

Command Default

All routes computed by IS-IS are installed in the RIB.

Command Modes

IS-IS router address-family

Command History

Release	Modification
Release 6.5.1	This command was introduced.

Usage Guidelines

When **distribute-list in** command is configured, some routes that IS-IS computes are not installed in the forwarding plane of the local router, but other IS-IS routers will not be aware of this. This introduces a difference between the forwarding state computed by other IS-IS routers and the actual forwarding state on this router. In some cases, this could lead to traffic being dropped or looped. Hence, be careful about when to use this command.

Task ID

Task ID	Operation
isis	read, write

Example

The following example shows how to filter the incoming IS-IS updates using prefix lists:

```
ipv6 prefix-list v6-1
 10 deny 2020:1000::101/128
 20 permit ::/0 le 128
!
ipv4 prefix-list v4-1
 10 deny 20.1.0.101/32
 20 permit 0.0.0.0/0 le 32
!
router isis 1
 net 49.0001.0000.0000.0100.00
```

```
 address-family ipv4 unicast
  metric-style wide
  distribute-list prefix-list v4-1 in
 !
 address-family ipv6 unicast
  metric-style wide
  distribute-list prefix-list v6-1 in
 !
```

The following example shows how to filter the incoming IS-IS updates using a route-policy for both IPv4 and IPv6:

```
prefix-set isis-ps-1
  2020:1000::101/128 le 128,
  20.1.0.101/32 le 32
end-set
!
route-policy isis-rp-1
  if destination in isis-ps-1 then
    drop
  elseif tag eq 0 then
    pass
  elseif tag le 100 then
    drop
  else
    pass
  endif
end-policy
!
router isis 1
 net 49.0001.0000.0000.0100.00
 address-family ipv4 unicast
  metric-style wide
  distribute-list route-policy isis-rp-1 in
 !
 address-family ipv6 unicast
  metric-style wide
  distribute-list route-policy isis-rp-1 in
 !
```

fast-reroute per-link (IS-IS)

To enable IP fast reroute (IPFRR) loop-free alternate (LFA) prefix independent per-link computation, use the **fast-reroute per-link** command in interface address family configuration mode. To disable this feature, use the **no** form of this command.

fast-reroute per-link [**exclude interface** *type interface-path-id* | **level** {**1** | **2**} | **lfa-candidate interface** *type interface-path-id*]
no fast-reroute per-link

Syntax Description

exclude	Specifies fast-reroute (FRR) loop-free alternate (LFA) computation exclusion information
level {1 | 2}	Configures FRR LFA computation for one level only.
lfa-candidate	Specifies FRR LFA computation candidate information
interface	Specifies an interface that needs to be either excluded from FRR LFA computation (when used with **exclude** keyword) or to be included to LFA candidate list in FRR LFA computation (when used with the **lfa-candidate** keyword).
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

IP fast-reroute LFA per-link computation is disabled.

Command Modes

Interface address family configuration

Command History

Release	Modification
Release 4.0.1	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
isis	read, write

This example shows how to configure per-link fast-reroute LFA computation for the IPv4 unicast topology at Level 1:

```
RP/0/RSP0/CPU0:router(config)# router isis isp
RP/0/RSP0/CPU0:router(config-isis)# interface GigabitEthernet 0/3/0/0
RP/0/RSP0/CPU0:router(config-isis-if)# address-family ipv4 unicast
RP/0/RSP0/CPU0:router(config-isis-if-af)# fast-reroute per-link level 1
```

fast-reroute per-prefix (IS-IS)

To enable IP fast reroute (IPFRR) loop-free alternate (LFA) prefix dependent computation, use the **fast-reroute per-prefix** command in interface address family configuration mode. LFA is supported only on Enhanced Ethernet line card.

fast-reroute per-prefix [**exclude interface** *type interface-path-id* | **level** {**1** | **2**} | **lfa-candidate interface** *type interface-path-id* | **remote-lfa** {**maximum-metric** *metric-value* | **tunnel mpls-ldp**} **prefix-list***prefix-list-name*[**level** {**1** | **2**}]]

Syntax Description

exclude	Specifies fast-reroute (FRR) loop-free alternate (LFA) computation exclusion information
level {1 | 2}	Configures FRR LFA computation for one level only.
lfa-candidate	Specifies FRR LFA computation candidate information
interface	Specifies an interface that needs to be either excluded from FRR LFA computation (when used with **exclude** keyword) or to be included to LFA candidate list in FRR LFA computation (when used with the **lfa-candidate** keyword).
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.
remote-lfa	Enable remote LFA related configuration.
prefix-list *prefix-list-name*	Filter PQ node router ID based on prefix list.

Command Default

IP fast-reroute LFA per-prefix computation is disabled.

Command Modes

Interface 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
isis	read, write

This example shows how to configure per-prefix fast-reroute LFA computation for the IPv4 unicast topology at Level 1:

```
RP/0/RSP0/CPU0:router(config)# router isis isp
RP/0/RSP0/CPU0:router(config-isis)# interface GigabitEthernet 0/3/0/0
RP/0/RSP0/CPU0:router(config-isis-if)# address-family ipv4 unicast
RP/0/RSP0/CPU0:router(config-isis-if-af)# fast-reroute per-prefix level 1
```

This example shows how to configure per-prefix remote-lfa prefix list. The prefix-list option filters PQ node router ID based on prefix list.

```
RP/0/RP0/CPU0:router(config-isis-af)# fast-reroute per-prefix remote-lfa prefix-list
```

fast-reroute per-link priority-limit (IS-IS)

To enable the IP fast reroute (IPFRR) loop-free alternate (LFA) prefix independent per-link computation, use the **fast-reroute per-link priority-limit** command in address family configuration mode. To disable this feature, use the **no** form of this command.

fast-reroute per-link priority-limit {**critical** | **high** | **medium**} **level** {**1** | **2**}
no fast-reroute per-link priority-limit

Syntax Description

critical	Enables LFA omputation for critical priority prefixes only.
high	Enables LFA computation for for criticaland high priority prefixes.
medium	Enables LFA computation for for critical, high, and medium priority prefixes.
level {1|2}	Sets priority-limit for routing Level 1 or Level 2 independently.

Command Default

Fast-reroute per link priority limit LFA computation is disabled.

Command Modes

IPv4 unicast address family configuration

IPv6 unicast address family configuration

IPv4 multicast address family configuration

IPv6 multicast address family configuration

Command History

Release	Modification
Release 4.0.1	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	Operations
isis	read, write

This example shows how to configure fast-reroute prefix independent per-link computation for critical priority prefixes for level 1 only:

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#router isis isp_lfa
RP/0/RSP0/CPU0:router(config-isis)#address-family ipv4
```

```
RP/0/RSP0/CPU0:router(config-isis-af)#fast-reroute per-link priority-limit critical level
1
```

fast-reroute per-link use-candidate-only (IS-IS)

To exclude all interfaces by default from prefix independent per-link fast-reroute computation, use the **fast-reroute per-link use-candidate-only** command ISIS address family configuration mode. To include all interfaces by default, use the no form of this command.

fast-reroute per-link use-candidate-only [**level** {**1** | **2**}]
no fast-reroute per-link use-candidate-only

Syntax Description

level {1 | 2}	[Optional] Sets priority-limit for routing Level 1 or Level 2 independently.

Command Default

None

Command Modes

IPv4 unicast address family configuration

IPv6 unicast address family configuration

IPv4 multicast address family configuration

IPv6 multicast address family configuration

Command History

Release	Modification
Release 4.2.0	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.

This command can be configured for a specific level. If no level is given it is applicable for both level 1 and level 2.

Task ID

Task ID	Operation
isis	read, write

This example shows how to configure the **fast-reroute per-link use-candidate-only** command for Level 1.

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#router isis isp
RP/0/RSP0/CPU0:router(config-isis)#address-family ipv4 unicast
RP/0/RSP0/CPU0:router(config-isis-af)#fast-reroute use-candidate-only per-link level 1
```

fast-reroute per-prefix load-sharing disable (IS-IS)

To disable load sharing prefixes across multiple backups, use the **fast-reroute per-prefix load-sharing disable** command in IPv4 address family configuration mode. To disable this feature, use the **no** form of this command.

fast-reroute per-prefix load-sharingdisable
no fast-reroute per-prefix load-sharingdisable

Syntax Description

level {1|2}	Disables load-sharing for Level 1 or Level 2 independently.

Command Default

Load sharing is enabled.

Command Modes

IPv4 unicast address family configuration

IPv4 multicast address family configuration

Command History

Release	Modification
Release 4.0.1	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	Operations
isis	read, write

This example shows how to disable load-sharing prefixes across multiple backups for level 1 routes:

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#router isis isp_lfa
RP/0/RSP0/CPU0:router(config-isis)#address-family ipv4
RP/0/RSP0/CPU0:router(config-isis-af)#fast-reroute per-prefix load-sharing disable level 1
```

fast-reroute per-prefix tiebreaker (IS-IS)

To configure tie-breaker for multiple backups, use the **fast-reroute per-prefix tiebreaker** command in IPv4 address family configuration mode. To disable tie-breaker configuration, use the **no** form of this command.

fast-reroute per-prefix tiebreaker [**downstream** | **lc-disjoint** | **lowest-backup-metric** | **node-protecting** | **primary-path** | **secondary-path** | **srlg-disjoint** *index*] **index** *index_number* **level {1 | 2}**
no fast-reroute per-prefix tiebreaker

Syntax Description

downstream	Configures to prefer backup path via downstream node, in case of tie-breaker.
lc-disjoint	Configures to prefer Prefer line card disjoint backup path.
lowest-backup-metric	Configures to prefer backup path with lowest total metric.
node-protecting	Configures to prefer node protecting backup path.
primary-path	Configures to prefer backup path from ECMP set.
secondary-path	Configures to prefer non-ECMP backup path.
srlg-disjoint	Configures to prefer srlg disjoint backup path.
index	Sets preference order among tie-breakers.
index_number	Value for the index. Range is 1-255.
level {1 | 2}	Configures tiebreaker for Level 1 or Level 2 independently.

Command Default

Tie-breaker for multiple backups is not configured.

Command Modes

IPv4 unicast address family configuration

IPv4 multicast address family configuration

Command History

Release	Modification
Release 4.0.1	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	Operations
isis	read, write

This example shows how to configure preference of backup path via downstream node in case of a tie-breaker for selection of backup path from multiple backup paths:

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#router isis isp_lfa
RP/0/RSP0/CPU0:router(config-isis)#address-family ipv4
RP/0/RSP0/CPU0:router(config-isis-af)#fast-reroute per-prefix tiebreaker downstream index
255
```

This example shows how to configure all the criterions for backup path selection:

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#router isis isp_lfa
RP/0/RSP0/CPU0:router(config-isis)#address-family ipv4
RP/0/RSP0/CPU0:router(config-isis-af)#fast-reroute per-prefix tiebreaker srlg-disjoint index
 10
RP/0/RSP0/CPU0:router(config-isis-af)#fast-reroute per-prefix tiebreaker primary-path index
 20
RP/0/RSP0/CPU0:router(config-isis-af)#fast-reroute per-prefix tiebreaker lowest-backup-metric
 index 30
RP/0/RSP0/CPU0:router(config-isis-af)#fast-reroute per-prefix tiebreaker lc-disjoint index
 40
RP/0/RSP0/CPU0:router(config-isis-af)#fast-reroute per-prefix tiebreaker node-protecting
index 50
```

fast-reroute per-prefix use-candidate-only (IS-IS)

To exclude all interfaces by default from prefix dependent fast-reroute computation, use the **fast-reroute per-prefix use-candidate-only** command ISIS address family configuration mode. To include all interfaces by default, use the **no** form of this command.

fast-reroute per-prefix use-candidate-only [**level** {**1** | **2**}]
no fast-reroute per-prefix use-candidate-only

Syntax Description

level {1 | 2}	[Optional] Sets priority-limit for routing Level 1 or Level 2 independently.

Command Default

None

Command Modes

IPv4 unicast address family configuration

IPv6 unicast address family configuration

IPv4 multicast address family configuration

IPv6 multicast address family configuration

Command History

Release	Modification
Release 4.2.0	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.

This command can be configured for a specific level. If no level is given it is applicable for both level 1 and level 2.

Task ID

Task ID	Operation
isis	read, write

This example shows how to configure the **fast-reroute per-prefix use-candidate-only** command for Level 1.

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#router isis isp
RP/0/RSP0/CPU0:router(config-isis)#address-family ipv4 unicast
RP/0/RSP0/CPU0:router(config-isis-af)#fast-reroute use-candidate-only per-link level 1
```

hello-interval (IS-IS)

To specify the length of time between consecutive hello packets sent by the Intermediate System-to-Intermediate System (IS-IS) protocol software, use the **hello-interval** command in interface configuration mode. To restore the default value, use the **no** form of this command.

hello-interval *seconds* [**level** {**1** | **2**}]
no hello-interval [*seconds*] [**level** {**1** | **2**}]

Syntax Description

seconds	Integer value (in seconds) for the length of time between consecutive hello packets. By default, a value three times the hello interval *seconds* is advertised as the *hold time* in the hello packets sent. (That multiplier of three can be changed by using the **hello-multiplier** command.) With smaller hello intervals, topological changes are detected more quickly, but there is more routing traffic. Range is 1 to 65535 seconds.
level { **1** | **2** }	(Optional) Specifies the hello interval for Level 1 and Level 2 independently. For broadcast interfaces only.

Command Default

seconds : 10 seconds

Both Level 1 and Level 2 are configured if no level is specified.

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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

The hello interval can be configured independently for Level 1 and Level 2, except on serial point-to-point interfaces. (Because only a single type of hello packet is sent on serial links, it is independent of Level 1 or Level 2.) Configuring Level 1 and Level 2 independently is used on LAN interfaces.

Note A shorter hello interval gives quicker convergence, but increases bandwidth and CPU usage. It might also add to instability in the network.

A slower hello interval saves bandwidth and CPU. Especially when used in combination with a higher hello multiplier, this strategy may increase overall network stability.

For point-to-point links, IS-IS sends only a single hello for Level 1 and Level 2, making the **level** keyword meaningless on point-to-point links. To modify hello parameters for a point-to-point interface, omit the **level** keyword.

Note Currently, a user can configure an aggressive hello-interval (lower than the default of 10 seconds for peer-to-peer session). But, if NSR or NSF is configured, the default hello interval has to be used so that the sessions do not run into the risk of flapping during switchover.

Using LAN adjacencies in high availability (HA) scenarios is not recommended, since there is no designated intermediate system (DIS) redundancy in the protocol and traffic will either drop or be rerouted temporarily during DIS re-election.

Task ID

Task ID	Operations
isis	read, write

Examples

The following example shows how to configure TenGigE interface 0/6/0/0 to advertise hello packets every 5 seconds for Level 1 topology routes. This situation causes more traffic than configuring a longer interval, but topological changes are detected more quickly.

```
RP/0/RSP0/CPU0:router(config)# router isis isp
RP/0/RSP0/CPU0:router(config-isis)# interface TenGigE 0/6/0/0
RP/0/RSP0/CPU0:router(config-isis-if)# hello-interval 5 level 1
```

Related Commands

Command	Description
hello-multiplier, on page 848	Specifies the number of IS-IS hello packets a neighbor must miss before the router should declare the adjacency as down.

hello-multiplier

To specify the number of Intermediate System-to-Intermediate System (IS-IS) hello packets a neighbor must miss before the router should declare the adjacency as down, use the **hello-multiplier** command in interface configuration mode. To restore the default value, use the **no** form of this command.

hello-multiplier *multiplier* [**level** {**1** | **2**}]
no hello-multiplier [*multiplier*] [**level** {**1** | **2**}]

Syntax Description

multiplier	Advertised hold time in IS-IS hello packets is set to the hello multiplier times the hello interval. Range is 3 to 1000. Neighbors declare an adjacency to this down router after not having received any IS-IS hello packets during the advertised hold time. The hold time (and thus the hello multiplier and the hello interval) can be set on an individual interface basis, and can be different between different networking devices in one area. Using a smaller hello multiplier gives faster convergence, but can result in more routing instability. Increase the hello multiplier to a larger value to help network stability when needed. Never configure a hello multiplier to a value lower than the default value of 3.
level { **1** | **2** }	(Optional) Specifies the hello multiplier independently for Level 1 or Level 2 adjacencies.

Command Default

multiplier : 3

Both Level 1 and Level 2 are configured if no level is specified.

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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

The “holding time” carried in an IS-IS hello packet determines how long a neighbor waits for another hello packet before declaring the neighbor to be down. This time determines how quickly a failed link or neighbor is detected so that routes can be recalculated.

Use the **hello-multiplier** command in circumstances where hello packets are lost frequently and IS-IS adjacencies are failing unnecessarily. You can raise the hello multiplier and lower the hello interval (hello-interval (IS-IS), on page 846 command) correspondingly to make the hello protocol more reliable without increasing the time required to detect a link failure.

On point-to-point links, there is only one hello for both Level 1 and Level 2. Separate Level 1 and Level 2 hello packets are also sent over nonbroadcast multiaccess (NBMA) networks in multipoint mode, such as X.25, Frame Relay, and ATM.

Task ID

Task ID	Operations
isis	read, write

Examples

The following example shows how the network administrator wants to increase network stability by making sure an adjacency goes down only when many (ten) hello packets are missed. The total time to detect link failure is 60 seconds. This strategy ensures that the network remains stable, even when the link is fully congested.

```
RP/0/RSP0/CPU0:router(config)# router isis isp
RP/0/RSP0/CPU0:router(config-isis)# interface GigabitEthernet /2/0/1
RP/0/RSP0/CPU0:router(config-isis-if)# hello-interval 6
RP/0/RSP0/CPU0:router(config-isis-if)# hello-multiplier 10
```

Related Commands

Command	Description
hello-interval (IS-IS), on page 846	Specifies the length of time between hello packets that the software sends.

hello-padding

To configure padding on Intermediate System-to-Intermediate System (IS-IS) hello protocol data units (IIH PDUs) for all IS-IS interfaces on the router, use the **hello-padding** command in interface configuration mode. To suppress padding, use the **no** form of this command.

hello-padding {**disable** | **sometimes**} [**level** {**1** | **2**}]
no hello-padding {**disable** | **sometimes**} [**level** {**1** | **2**}]

Syntax Description

disable	Suppresses hello padding.
sometimes	Enables hello padding during adjacency formation only.
level { **1** | **2** }	(Optional) Specifies hello padding for Level 1 or Level 2 independently.

Command Default

Hello padding is enabled.

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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

You might want to suppress hello padding to conserve network resources. The lower the circuit speed, the higher the percentage of padding overhead. Before suppressing the hello padding, you should know your physical and data link layer configurations and have control over them, and also know your router configuration at the network layer.

For point-to-point links, IS-IS sends only a single hello for Level 1 and Level 2, making the **level** keyword meaningless on point-to-point links. To modify hello parameters for a point-to-point interface, omit the **level** keyword.

Task ID

Task ID	Operations
isis	read, write

Examples

The following example shows how to suppress IS-IS hello padding over local area network (LAN) circuits for interface GigabitEthernet 0/2/0/1:

```
RP/0/RSP0/CPU0:router(config)# router isis isp
RP/0/RSP0/CPU0:router(config-isis)# interface GigabitEthernet 0/2/0/1
RP/0/RSP0/CPU0:router(config-isis-if)# hello-padding disable
```

Related Commands

Command	Description
show isis interface, on page 979	Displays information about the IS-IS interface.

hello-password

To configure the authentication password for an Intermediate System-to-Intermediate System (IS-IS) interface, use the **hello-password** command in interface configuration mode. To disable authentication, use the **no** form of this command.

hello-password [**hmac-md5** | **text**] [**clear** | **encrypted**] *password* [**level** {**1** | **2**}] [**send-only**]
no hello-password [**hmac-md5** | **text**] [**clear** | **encrypted**] *password* [**level** {**1** | **2**}] [**send-only**]

Syntax Description

hmac-md5	(Optional) Specifies that the password use HMAC-MD5 authentication.
text	(Optional) Specifies that the password use clear text password authentication.
clear	(Optional) Specifies that the password be unencrypted.
encrypted	(Optional) Specifies that the password be encrypted using a two-way algorithm.
password	Authentication password you assign for an interface.
level { **1** | **2** }	(Optional) Specifies whether the password is for a Level 1 or a Level 2 protocol data unit (PDU).
send-only	(Optional) Specifies that the password applies only to protocol data units (PDUs) that are being sent and does not apply to PDUs that are being received.

Command Default

Both Level 1 and Level 2 are configured if no level is specified.

password: encrypted text

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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

When a **text** password is configured, it is exchanged as clear text. Therefore, the **hello-password** command provides limited security.

When an **hmac-md5** password is configured, the password is never sent over the network and is instead used to calculate a cryptographic checksum to ensure the integrity of the exchanged data.

For point-to-point links, IS-IS sends only a single hello for Level 1 and Level 2, making the **level** keyword meaningless on point-to-point links. To modify hello parameters for a point-to-point interface, omit the **level** keyword.

Task ID

Task ID	Operations
isis	read, write

Examples

The following example shows how to configure a password with HMAC-MD5 authentication for hello packets running on GigabitEthernet 0/2/0/3 interface:

```
RP/0/RSP0/CPU0:router(config)# router isis isp
RP/0/RSP0/CPU0:router(config-isis)# interface GigabitEthernet 0/2/0/3
RP/0/RSP0/CPU0:router(config-isis-if)# hello-password hmac-md5 clear mypassword
```

Related Commands

Command	Description
hello-password keychain, on page 854	Configures the authentication password keychain for an Intermediate System-to-Intermediate System (IS-IS) interface.
hello-password accept, on page 856	Configures an additional authentication password for an IS-IS interface.

hello-password keychain

To configure the authentication password keychain for an Intermediate System-to-Intermediate System (IS-IS) interface, use the **hello-password keychain** command in interface configuration mode. To disable the authentication password keychain, use the **no** form of this command.

hello-password keychain *keychain-name* [**level** {**1** | **2**}] [**send-only**]
no hello-password keychain *keychain-name* [**level** {**1** | **2**}] [**send-only**]

Syntax Description

keychain	Keyword that specifies the keychain to be configured. An authentication password keychain is a sequence of keys that are collectively managed and used for authenticating a peer-to-peer group.
keychain-name	Specifies the name of the keychain.
level { **1** | **2** }	(Optional) Specifies whether the keychain is for a Level 1 or a Level 2 protocol data unit (PDU).
send-only	(Optional) Specifies that the keychain applies only to protocol data units (PDUs) that are being sent and does not apply to PDUs that are being received.

Command Default

Both Level 1 and Level 2 are configured if no level is specified.

password: encrypted text

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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Specify a keychain to enable keychain authentication between two IS-IS peers. Use the **keychain** *keychain-name* keyword and argument to implement hitless key rollover for authentication.

Task ID

Task ID	Operations
isis	read, write

Examples

The following example shows how to configure a password keychain for level 1, send only authentication on a GigabitEthernet interface:

```
RP/0/RSP0/CPU0:routerRP/0/RSP0/CPU0:router(config)# router isis isp
RP/0/RSP0/CPU0:routerRP/0/RSP0/CPU0:router(config-isis)# interface GigabitEthernet 0/1/0/0
```

```
RP/0/RSP0/CPU0:routerRP/0/RSP0/CPU0:router(config-isis-if)# hello-password keychain
mykeychain level 1 send-only
```

Related Commands

Command	Description
hello-password, on page 852	Configures the authentication password for an Intermediate System-to-Intermediate System (IS-IS) interface.
hello-password accept, on page 856	Configures an additional authentication password for an IS-IS interface.

hello-password accept

To configure an additional authentication password for an Intermediate System-to-Intermediate System (IS-IS) interface, use the **hello-password accept** command in interface configuration mode. To disable authentication, use the **no** form of this command.

hello-password accept {**clear** | **encrypted**} *password* [**level** {**1** | **2**}]
no hello-password accept {**clear** | **encrypted**} *password* [**level** {**1** | **2**}]

Syntax Description

clear	Specifies that the password be unencrypted.
encrypted	Specifies that the password be encrypted using a two-way algorithm.
password	Authentication password you assign.
level { **1** | **2** }	(Optional) Specifies the password for Level 1 or Level 2 independently.

Command Default

Both Level 1 and Level 2 are configured if no level is specified.

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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **hello-password accept** command to add an additional password for an IS-IS interface. An authentication password must be configured using the **hello-password** command before an accept password can be configured for the corresponding level.

Task ID

Task ID	Operations
isis	read, write

Examples

The following example shows how to configure a password:

```
RP/0/RSP0/CPU0:routerRP/0/RSP0/CPU0:router(config)# router isis isp
RP/0/RSP0/CPU0:routerRP/0/RSP0/CPU0:router(config-isis)# interface GigabitEthernet 0/2/0/3

RP/0/RSP0/CPU0:routerRP/0/RSP0/CPU0:router(config-isis)# hello-password accept encrypted
111D1C1603
```

Related Commands

Command	Description
hello-password, on page 852	Configures an authentication password for an IS-IS interface.

hostname dynamic disable

To disable Intermediate System-to-Intermediate System (IS-IS) routing protocol dynamic hostname mapping, use the **hostname dynamic** command in router configuration mode. To remove the specified command from the configuration file and restore the system to its default condition, use the **no** form of this command.

hostname dynamic disable
no hostname dynamic disable

Syntax Description

disable	Disables dynamic host naming.

Command Default

Router names are dynamically mapped to system IDs.

Command Modes

Router configuration

Command History

Release	Modification
Release 3.7.2	This command was introduced.

Usage Guidelines

In an IS-IS routing domain, each router is represented by a 6-byte hexadecimal system ID. When network administrators maintain and troubleshoot networking devices, they must know the router name and corresponding system ID.

Link-state packets (LSPs) include the dynamic hostname in the type, length, and value (TLV) which carries the mapping information across the entire domain. Every router in the network, upon receiving the TLV from an LSP, tries to install it in a mapping table. The router then uses the mapping table when it wants to convert a system ID to a router name.

To display the entries in the mapping tables, use the **show isis hostname** command.

Task ID

Task ID	Operations
isis	read, write

Examples

The following example shows how to disable dynamic mapping of hostnames to system IDs:

```
RP/0/RSP0/CPU0:router(config)# router isis isp
 RP/0/RSP0/CPU0:router(config-isis)# hostname dynamic disable
```

Related Commands

Command	Description
hostname	Specifies the name of the local router.
show isis hostname, on page 977	Displays the router name-to-system ID mapping table.

ignore-lsp-errors

To override the default setting of a router to ignore Intermediate System-to-Intermediate System (IS-IS) link-state packets (LSPs) that are received with internal checksum errors, use the **ignore-lsp-errors disable** command in router configuration mode. To enable ignoring IS-IS LSP errors, use the **no** form of this command.

ignore-lsp-errors disable
no ignore-lsp-errors disable

Syntax Description

disable	Disables the functionality of the command.

Command Default

The system ignores corrupt LSPs.

Command Modes

Router 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

The IS-IS protocol definition requires that a received LSP with an incorrect data-link checksum be purged by the receiver, which causes the initiator of the packet to regenerate it. However, if a network has a link that causes data corruption and at the same time is delivering LSPs with correct data-link checksums, a continuous cycle of purging and regenerating large numbers of packets can occur. Because this situation could render the network nonfunctional, use this command to ignore these LSPs rather than purge the packets.

The receiving network devices use link-state packets to maintain their routing tables.

Task ID

Task ID	Operations
isis	read, write

Examples

The following example shows how to instruct the router to purge LSPs that cause the initiator to regenerate LSPs:

```
RP/0/RSP0/CPU0:router(config)# router isis isp
RP/0/RSP0/CPU0:router(config-isis)# ignore-lsp-errors disable
```

instance-id

To allow a router to share one or more circuits among multiple Intermediate System to Intermediate System (IS-IS) routing protocol instances, use the **instance-id** command in router configuration mode.

instance-id *identifier*

Syntax Description

identifier	Specifies the Intermediate System to Intermediate System (IS-IS) routing protocol instance. Range is 1-65535.

Command Default

Disabled

Command Modes

Router configuration

Command History

Release	Modification
Release 6.1.x	This command was introduced.

Task ID

Task ID	Operations
isis	read, write

Examples

The following example shows how to configure multiple instances on a single router:

```
RP/0/RSP0/CPU0:router(config)# router isis ring
RP/0/RSP0/CPU0:router(config-isis)# instance-id 1
RP/0/RSP0/CPU0:router(config-isis)# exit
RP/0/RSP0/CPU0:router(config)# router isis 1
RP/0/RSP0/CPU0:router(config-isis)# instance-id 6
RP/0/RSP0/CPU0:router(config-isis)#
```

interface (IS-IS)

To configure the Intermediate System-to-Intermediate System (IS-IS) protocol on an interface, use the **interface** command in router configuration mode. To disable IS-IS routing for interfaces, use the **no** form of this command.

interface *type interface-path-id*
no interface *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

No interfaces are specified.

Command Modes

Router 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

An address family must be established on the IS-IS interface before the interface is enabled for IS-IS protocol operation.

Task ID

Task ID	Operations
isis	read, write

Examples

The following example shows how to enable an IS-IS multitopology configuration for IPv4 on GigabitEthernet interface 0/3/0/0:

```
RP/0/RSP0/CPU0:router(config)# router isis isp
RP/0/RSP0/CPU0:router(config-isis)# net 49.0000.0000.0001.00
RP/0/RSP0/CPU0:router(config-isis)# interface GigabitEthernet 0/3/0/0
RP/0/RSP0/CPU0:router(config-isis-if)# address-family ipv4 unicast
RP/0/RSP0/CPU0:router(config-isis-if-af)# metric-style wide level 1
```

```
!
RP/0/RSP0/CPU0:router(config)# interface GigabitEthernet 0/3/0/0
RP/0/RSP0/CPU0:router(config-if)# ipv4 address 2001::1/64
```

Related Commands

Command	Description
log adjacency changes (IS-IS), on page 872	Configures the routing level for an instance of the IS-IS routing process.
net, on page 922	Configures an IS-IS network entity title (NET) for the routing process.
router isis, on page 950	Enables the IS-IS routing protocol.

ipfrr lfa

The ipfrr lfa commands are deprecated in Release 4.0.1, and replaced with the fast-reroute commands. For more information, see the following:

- fast-reroute per-link priority-limit [level <1-2>]
- fast-reroute per-prefix priority-limit [level <1-2>]
- fast-reroute per-prefix tiebreaker index <1-255> [level <1-2>]
- fast-reroute per-prefix load-sharing disable [level <1-2>]
- fast-reroute [level <1-2>]
- fast-reroute per-link exclude interface [level <1-2>]
- fast-reroute per-link lfa-candidate interface [level <1-2>]
- fast-reroute per-prefix exclude interface [level <1-2>]
- fast-reroute per-prefix lfa-candidate interface [level <1-2>]
- show isis fast-reroute summary
- show isis fast-reroute [prefix] [longer-prefixes]
- show isis fast-reroute detail [prefix] [longer-prefixes]

To enable the IP fast reroute (IPFRR) loop-free alternate (LFA) computation, use the **ipfrr lfa** command in interface address family configuration mode. To disable this feature, use the **no** form of this command.

ipfrr lfa level {**1** | **2**}
no ipfrr lfa level {**1** | **2**}

Syntax Description

level { **1** | **2** }	Configures IPFRR LFA for Level 1 or Level 2 independently.

Command Default

IPFRR LFA is disabled.

Command Modes

Interface address family configuration

Command History

Release	Modification
Release 3.9.0	This command was introduced.
Release 4.0.1	This command was deprecated and replaced with the fast-reroute commands.

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.

Use the **ipfrr lfa** command to compute loop-free alternates for all links or neighbors in the event of a link failure.

To enable node protection on broadcast links, IPRR and bidirectional forwarding detection (BFD) must be enabled on the interface under IS-IS. See *Cisco IOS XR Interface and Hardware Configuration Guide* for information on configuring BFD.



Note Multiprotocol Label Switching (MPLS) FRR and IPFRR cannot be configured on the same interface simultaneously.

Task ID

Task ID	Operations
isis	read, write

Examples

The following example shows how to configure IPFRR for the IPv4 unicast topology at Level 1:

```
RP/0/RSP0/CPU0:router(config)# router isis isp
RP/0/RSP0/CPU0:router(config-isis)# interface GigabitEthernet 0/3/0/0
RP/0/RSP0/CPU0:router(config-isis-if)# address-family ipv4 unicast
RP/0/RSP0/CPU0:router(config-isis-if-af)# ipfrr lfa level 1
```

ipfrr lfa exclude interface

To exclude an interface from the IP fast reroute (IPFRR) loop-free alternate (LFA) computation, use the **ipfrr lfa exclude interface** command in interface address family configuration mode. To disable this feature, use the **no** form of this command.

ipfrr lfa exclude interface *type interface-path-id*
no ipfrr lfa exclude interface *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

IPFRR LFA is disabled.

Command Modes

Interface address family configuration

Command History

Release	Modification
Release 3.9.0	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.

Use the **ipfrr lfa** command to compute loop-free alternates for all links or neighbors in the event of a link failure.

To enable node protection on broadcast links, IPRR and bidirectional forwarding detection (BFD) must be enabled on the interface under IS-IS. See *Cisco IOS XR Interface and Hardware Configuration Guide* for information on configuring BFD.

Note Multiprotocol Label Switching (MPLS) FRR and IPFRR cannot be configured on the same interface simultaneously.

Task ID

Task ID	Operations
isis	read, write

Examples

The following example shows how to configure to exclude 0/1/0/0 interface from IPFRR LFA:

```
RP/0/RSP0/CPU0:router(config)# router isis isp
RP/0/RSP0/CPU0:router(config-isis)# interface GigabitEthernet 0/1/0/0
RP/0/RSP0/CPU0:router(config-isis-if)# address-family ipv4 unicast
RP/0/RSP0/CPU0:router(config-isis-if-af)# ipfrr lfa exclude interface GigabitEthernet 0/1/0/0
```

Related Commands

Command	Description
ipfrr lfa, on page 863	Enable the IP fast reroute (IPFRR) loop-free alternate (LFA) computation

ispf

To configure the incremental shortest path first (iSPF) algorithm to calculate network topology, use the **ispf** command in address family configuration mode. To disable this algorithm function, use the **no** form of this command.

ispf [**level** {**1** | **2**}]
no ispf [**level** {**1** | **2**}]

Syntax Description

level { **1** | **2** }	(Optional) Configures the iSPF algorithm for Level 1 or Level 2 independently.

Command Default

The iSPF algorithm is not configured.

Command Modes

Address family 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

The iSPF algorithm may be used to reduce the processor load when IS-IS needs to recalculate its topology after minor changes.



Note The incremental shortest path first (ISPF) algorithm was developed to improve SPF performance in the days of slow CPUs and before more advanced technological solutions like Fast Re-Route (FRR) and micro-loop avoidance were available. FRR and micro-loop avoidance both depend on information that is computed during a full SPF calculation and saved in the resulting shortest path tree.

Task ID

Task ID	Operations
isis	read, write

Examples

The following example shows how to configure iSPF for the IPv4 unicast topology at Level 1:

```
RP/0/RSP0/CPU0:router(config)# router isis isp
RP/0/RSP0/CPU0:router(config-isis)# address-family ipv4 unicast
RP/0/RSP0/CPU0:router(config-isis-af)# ispf level 1
```

is-type

To configure the routing level for an Intermediate System-to-Intermediate System (IS-IS) area, use the **is-type** command in router configuration mode. To set the routing level to the default level, use the **no** form of this command.

is-type {**level-1** | **level-1-2** | **level-2-only**}
no is-type [**level-1** | **level-1-2** | **level-2-only**]

Syntax Description

level-1	Specifies that the router perform only Level 1 (intra-area) routing. This router learns only about destinations inside its area. Level 2 (interarea) routing is performed by the closest Level 1-2 router.
level-1-2	Specifies that the router perform both Level 1 and Level 2 routing.
level-2-only	Specifies that the routing process acts as a Level 2 (interarea) router only. This router is part of the backbone, and does not communicate with Level 1-only routers in its own area.

Command Default

Both Level 1 and Level 2 are configured if no level is specified.

Command Modes

Router 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

When the router is configured with Level 1 routing only, this router learns about destinations only inside its area. Level 2 (interarea) routing is performed by the closest Level 1-2 router.

When the router is configured with Level 2 routing only, this router is part of the backbone, and does not communicate with Level 1 routers in its own area.

The router has one link-state packet database (LSDB) for destinations inside the area (Level 1 routing) and runs a shortest path first (SPF) calculation to discover the area topology. It also has another LSDB with link-state packets (LSPs) of all other backbone (Level 2) routers, and runs another SPF calculation to discover the topology of the backbone and the existence of all other areas.

We highly recommend that you configure the type of an IS-IS routing process to establish the proper level of adjacencies. If there is only one area in the network, there is no need to run both Level 1 and Level 2 routing algorithms.

Task ID

Task ID	Operations
isis	read, write

Examples

The following example shows how to specify that the router is part of the backbone and that it does not communicate with Level 1-only routers:

```
RP/0/RSP0/CPU0:router(config)# router isis isp
RP/0/RSP0/CPU0:router(config-isis)# is-type level-2-only
```

Related Commands

Command	Description
circuit-type, on page 817	Configures the type of adjacency.
show isis neighbors, on page 992	Displays information about IS-IS neighbors.

link-down fast-detect

To enable fast detection of link down events under ISIS, use the **link-down fast-detect** command in ISIS interface configuration mode.

link-down fast-detect

Syntax Description

This command has no arguments or keywords.

Command Default Fast detect is not enabled.

Command Modes ISIS interface configuration.

Command History

Release	Modification
Release 3.9.1	This command was introduced.

Example

The following example shows how to enable fast detection of link down events under ISIS:

```
RP/0/RP0/CPU0:router(config)# router isis isp
RP/0/RP0/CPU0:router(config-isis)# interface HundredGigE0/3/0/2
RP/0/RP0/CPU0:router(config-isis-if)# link-down fast-detect
```

link-group

To configure an interface as a member of a link group, use the **link-group** command in the IS-IS interface or address-family configuration mode. To remove an interface from a link-group, use the **no** form of this command.

link-group *link-group-name*
no link-group *link-group-name*

Syntax Description

link-group-name	Name of a link group.

Command Default

No link groups are configured.

Command Modes

IS-IS interface configuration

Address-family configuration

Command History

Release	Modification
Release 4.3.1	This command was introduced.

Usage Guidelines

One IS-IS interface and address-family can specify only one link-group association. The default is for both levels regardless of current circuit-type. The link-group association can be specified for one level only if configured.

Task ID

Task ID	Operations
isis	read, write

Examples

The following example shows how to configure GigabitEthernet interface 0/3/0/0 as a member of a link group:

```
RP/0/RSP0/CPU0:router(config)# router isis isp
RP/0/RSP0/CPU0:router(config-isis)# interface GigabitEthernet 0/3/0/0
RP/0/RSP0/CPU0:router(config-isis-if)# address-family ipv4 unicast
RP/0/RSP0/CPU0:router(config-isis-if-af)# link-group purple
```

log adjacency changes (IS-IS)

To cause an IS-IS instance to generate a log message when an Intermediate System-to-Intermediate System (IS-IS) adjacency changes state (up or down), use the **log adjacency changes** command in router configuration mode. To restore the default value, use the **no** form of this command.

log adjacency changes
no log adjacency changes

Command Default

No IS-IS instance log messages are generated.

Command Modes

Router configuration

Command History

Release	Modification
Release 3.7.2	This command was introduced.

Usage Guidelines

Use the **log adjacency changes** command to monitor IS-IS adjacency state changes; it may be very useful when you are monitoring large networks. Messages are logged using the system error message facility. Messages can be in either of two forms:

```
%ISIS-4-ADJCHANGE: Adjacency to 0001.0000.0008 (Gi 0/2/1/0) (L2) Up, new adjacency
%ISIS-4-ADJCHANGE: Adjacency to router-gsr8 (Gi /2/1/0) (L1) Down, Holdtime expired
```

Using the **no** form of the command removes the specified command from the configuration file and restores the system to its default condition with respect to the command.

Task ID

Task ID	Operations
isis	read, write

Examples

The following example shows how to configure the router to log adjacency changes:

```
RP/0/RSP0/CPU0:router(config)# router isis isp
RP/0/RSP0/CPU0:router(config-isis)# log adjacency changes
```

Related Commands

Command	Description
logging	Logs messages to a syslog server host.

log pdu drops

To log Intermediate System-to-Intermediate System (IS-IS) protocol data units (PDUs) that are dropped, use the **log pdu drops** command in router configuration mode. To disable this function, use the **no** form of this command.

log pdu drops
no log pdu drops

Command Default

PDU logging is disabled.

Command Modes

Router 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **log pdu drops** command to monitor a network when IS-IS PDUs are suspected of being dropped. The reason for the PDU being dropped and current PDU drop statistics are recorded.

The following are examples of PDU logging output:

```
%ISIS-4-ERR_LSP_INPUT_Q_OVERFLOW - An incoming LSP or SNP pdu was dropped because the input
 queue was full
%ISIS-3-ERR_SEND_PAK - The process encountered a software-error while sending the IS-IS
packet
```

Task ID

Task ID	Operations
isis	read, write

Examples

The following example shows how to enable PDU logging:

```
RP/0/RSP0/CPU0:router(config)# router isis isp
RP/0/RSP0/CPU0:router(config-isis)# log pdu drops
```

lsp fast-flood threshold

To configure the link-state packet (LSP) fast-flood threshold, use the **lsp fast-flood threshold** command in interface configuration mode. To restore the default value, use the **no** form of this command.

lsp fast-flood threshold *lsp-number* [**level** {**1** | **2**}]
no lsp fast-flood threshold [*lsp-number*] [**level** {**1** | **2**}]

Syntax Description

lsp-number	Number of LSPs to send back to back. Range is 1 to 4294967295.
level { **1** | **2** }	(Optional) Specifies the LSP threshold for Level 1 or Level 2 independently.

Command Default

10 LSPs are allowed in a back-to-back window

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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **lsp fast-flood threshold** command to accelerate convergence of LSP database. LSPs are sent back-to-back over an interface up to the specified limit. Past the limit, LSPs are sent out in the next batch window as determined by LSP pacing interval.

Duration of back-to-back window = LSP interval * LSP fast-flood threshold limit.

Task ID

Task ID	Operations
isis	read, write

Examples

The following example shows how to configure the LSP threshold:

```
RP/0/RSP0/CPU0:router(config)# router isis isp
RP/0/RSP0/CPU0:router(config-isis)# interface GigabitEthernet 0/3/0/0
RP/0/RSP0/CPU0:router(config-isis-if)# lsp fast-flood threshold 234 level 1
```

Related Commands

Command	Description
lsp-interval, on page 879	Configures the amount of time between consecutive LSPs sent on an IS-IS interface.

lsp-fast-flooding

To enable the fast flooding of Link State Packets (LSPs) to improve network response to topology changes, use the **lsp-fast-flooding** command in IS-IS router configuration mode.

lsp-fast-flooding [**max-lsp-tx** *lsps-per-sec* | **remote-psnp-delay** *milliseconds*\

Syntax Description

max-lsp-tx *lsps-per-sec*	(Optional) Sets the maximum number of LSPs that can be transmitted per second. The default is 1000. The range is from 33 through 5000
remote-psnp-delay *milliseconds*	(Optional) Specifies the maximum delay, in milliseconds, that the router should wait for a PSNP acknowledgment from neighbors after sending an LSP. The default is 500 ms. The range is from 1 through 5000 ms.

Command Default

This command is disabled by default. IS-IS fast flooding of LSPs is not enabled.

Command Modes

Router configuration

Command History

Release	Modification
Release 24.3.1	This command was introduced.

Usage Guidelines

None.

Task ID

Task ID	Operation
isis	read, write

Example

This example shows how to enable IS-IS fast flooding of LSPs.

```
Router# configure terminal
Router(config)# router isis 1
Router(config-router)# lsp-fast-flooding
```

This example shows how to set the maximum LSP transmission rate to 600 LSPs per second.

```
Router# configure terminal
Router(config)# router isis 1
Router(config-router)# lsp-fast-flooding
Router(config-isis-lsp-fast-flooding)# max-lsp-tx 600
```

This example shows how to set the PSNP acknowledgment delay to 600 milliseconds.

```
Router# configure terminal
Router(config)# router isis 1
Router(config-router)# lsp-fast-flooding
Router(config-isis-lsp-fast-flooding)# remote-psnp-delay 600
```

lsp-gen-interval

To customize IS-IS throttling of link-state packet (LSP) generation, use the **lsp-gen-interval** command in router configuration mode. To restore the default value, use the **no** form of this command.

lsp-gen-interval [**initial-wait** *initial*] [**secondary-wait** *secondary*] [**maximum-wait** *maximum*] [**level** {**1** | **2**}]
no lsp-gen-interval [[**initial-wait** *initial*] [**secondary-wait** *secondary*] [**maximum-wait** *maximum*]] [**level** {**1** | **2**}]

Syntax Description

initial-wait *initial*	Specifies the initial LSP generation delay (in milliseconds). Range is 0 to 120000 milliseconds.
secondary-wait *secondary*	Specifies the hold time between the first and second LSP generation (in milliseconds). Range is 1 to 120000 milliseconds.
maximum-wait *maximum*	Specifies the maximum interval (in milliseconds) between two consecutive occurrences of an LSP being generated. Range is 1 to 120000 milliseconds.
level { **1** | **2** }	(Optional) Specifies the LSP time interval for Level 1 or Level 2 independently.

Command Default

initial-wait *initial* : 50 milliseconds

secondary-wait *secondary* : 200 milliseconds

maximum-wait *maximum* : 5000 milliseconds

Command Modes

Router 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

During prolonged periods of network instability, repeated recalculation of LSPs can cause increased CPU load on the local router. Further, the flooding of these recalculated LSPs to the other Intermediate Systems in the network causes increased traffic and can result in other routers having to spend more time running route calculations.

Use the **lsp-gen-interval** command to reduce the rate of LSP generation during periods of instability in the network. This command can help to reduce CPU load on the router and to reduce the number of LSP transmissions to its IS-IS neighbors.

Task ID

Task ID	Operations
isis	read, write

Examples

The following example shows how to set the maximum interval between two consecutive occurrences of an LSP to 15 milliseconds and the initial LSP generation delta to 5 milliseconds:

```
RP/0/RSP0/CPU0:router(config)# router isis isp
RP/0/RSP0/CPU0:router(config-isis)# lsp-gen-interval maximum-wait 15 initial-wait 5
```

Related Commands

Command	Description
retransmit-interval (IS-IS), on page 946	Configures the amount of time between retransmission of each IS-IS LSP on a point-to-point link.

lsp-interval

To configure the amount of time between consecutive link-state packets (LSPs) sent on an Intermediate System-to-Intermediate System (IS-IS) interface, use the **lsp-interval** command in interface configuration mode. To restore the default value, use the **no** form of this command.

lsp-interval *milliseconds* [**level** {**1** | **2**}]
no lsp-interval [*milliseconds*] [**level** {**1** | **2**}]

Syntax Description

milliseconds	Time delay (in milliseconds) between successive LSPs. Range is 1 to 4294967295.
level { **1** | **2** }	(Optional) Configures the LSP time delay for Level 1 or Level 2 independently.

Command Default

milliseconds : 33 milliseconds

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 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
isis	read, write

Examples

The following example shows how to cause the system to send LSPs every 100 milliseconds (10 packets per second) on Level 1 and Level 2:

```
RP/0/RSP0/CPU0:router(config)# router isis isp
RP/0/RSP0/CPU0:router(config-isis)# interface GigabitEthernet /2/0/1
RP/0/RSP0/CPU0:router(config-isis-if)# lsp-interval 100
```

Related Commands

Command	Description
retransmit-interval (IS-IS), on page 946	Configures the amount of time between retransmission of each IS-IS LSP on a point-to-point link.

lsp-mtu

To set the maximum transmission unit (MTU) size of Intermediate System-to-Intermediate System (IS-IS) link-state packets (LSPs), use the **lsp-mtu** command in router configuration mode. To restore the default, use the **no** form of this command.

lsp-mtu *bytes* [**level** {**1** | **2**}]
no lsp-mtu [*bytes*] [**level** {**1** | **2**}]

Syntax Description

bytes	Maximum packet size in bytes. The number of bytes must be less than or equal to the smallest MTU of any link in the network. Range is 128 to 4352 bytes. **Note** Range is 128 to 8979 bytes from Release 6.6.3 onwards.
level { **1** | **2** }	(Optional) Specifies routing Level 1 or Level 2 independently.

Command Default

Both Level 1 and Level 2 are configured if no level is specified.

Command Modes

Router configuration

Command History

Release	Modification
Release 3.7.2	This command was introduced.
Release 3.9.1	The maximum packet size is increased to 8979.

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.

Under normal conditions, the default MTU size should be sufficient. However, if the MTU size of a link is less than 1500 bytes, the LSP MTU size must be lowered accordingly on each router in the network. If this action is not taken, routing becomes unpredictable.

This guideline applies to all Cisco networking devices in a network. If any link in the network has a reduced MTU size, all devices must be changed, not just the devices directly connected to the link.

Note Do not set the **lsp-mtu** command (network layer) to a value greater than the link MTU size that is set with the **mtu** command (physical layer).

To be certain about a link MTU size, use the show isis interface, on page 979 command to display the value.

Task ID

Task ID	Operations
isis	read, write

Examples

The following example shows how to set the MTU size to 1300 bytes:

```
RP/0/RSP0/CPU0:router(config)# router isis isp
RP/0/RSP0/CPU0:router(config-isis)# lsp-mtu 1300
```

Related Commands

Command	Description
mtu	Adjusts the maximum packet size or MTU size.
show isis interface, on page 979	Displays information about the IS-IS interface.

lsp-password

To configure the link-state packet (LSP) authentication password, use the **lsp-password** command in router configuration mode. To remove the **lsp-password** command from the configuration file and disable link-state packet authentication, use the **no** form of this command.

lsp-password [[**hmac-md5** | **text**] [**clear** | **encrypted**] *password* | **keychain** *keychain-name*] [**level** {**1** | **2**}] [**send-only**] [**snp send-only**]]
no lsp-password [[**hmac-md5** | **text**] [**clear** | **encrypted**] *password* | **keychain** *keychain-name*] [**level** {**1** | **2**}] [**send-only**] [**snp send-only**]]

Syntax Description

hmac-md5	Specifies that the password uses HMAC-MD5 authentication.
text	Specifies that the password uses clear text password authentication.
clear	Specifies that the password be unencrypted.
encrypted	Specifies that the password be encrypted using a two-way algorithm.
password	Authentication password you assign.
keychain	(Optional) Specifies a keychain.
keychain-name	Name of the keychain.
level { **1** | **2** }	(Optional) Specifies the password for Level 1 or Level 2 independently.
send-only	(Optional) Adds passwords to LSP and sequence number protocol (SNP) data units when they are sent. Does not check for authentication in received LSPs or sequence number PDUs (SNPs).
snp send-only	(Optional) Adds passwords to SNP data units when they are sent. Does not check for authentication in received SNPs. This option is available when the **text** keyword is specified.

Command Default

Both Level 1 and Level 2 are configured if no level is specified.

Command Modes

Router configuration

Command History

Release	Modification
Release 3.7.2	This command was introduced.

Usage Guidelines

When a **text** password is configured, it is exchanged as clear text. Therefore, the **lsp-password** command provides limited security.

When an **HMAC-MD5** password is configured, the password is never sent over the network and is instead used to calculate a cryptographic checksum to ensure the integrity of the exchanged data.

The recommended password configuration is that both incoming and outgoing SNPs be authenticated.

Note To disable SNP password checking, the **snp send-only** keywords must be specified in the **lsp-password** command.

To configure an additional password, use the **lsp-password accept** command.

Specify a key chain to enable key chain authentication between two IS-IS peers. Use the **keychain** *keychain-name* keyword and argument to implement hitless key rollover for authentication.

If you are performing LSP authentication and want to use the Purge Originator Identification feature, then use the **enable-poi** keyword in the **lsp-password** command.

Task ID

Task ID	Operations
isis	read, write

Examples

The following example shows how to configure separate Level 1 and Level 2 LSP and SNP passwords, one with HMAC-MD5 authentication and encryption and one with clear text password authentication and no encryption:

```
RP/0/RSP0/CPU0:router(config)# router isis isp
RP/0/RSP0/CPU0:router(config-isis)# lsp-password hmac-md5 clear password1 level 1
RP/0/RSP0/CPU0:router(config-isis)# lsp-password text clear password2 level 2
```

Related Commands

Command	Description
lsp-password accept, on page 884	Configures an additional LSP password when one LSP password is already configured for a level.

lsp-password accept

To configure an additional link-state packet (LSP) authentication password, use the **lsp-password accept** command in router configuration mode. To remove the **lsp-password accept** command from the configuration file and restore the system to its default condition, use the **no** form of this command.

lsp-password accept {**clear** | **encrypted**} *password* [**level** {**1** | **2**}]
no lsp-password accept [{**clear** | **encrypted**} *password* [**level** {**1** | **2**}]]

Syntax Description

clear	Specifies that the password be unencrypted.
encrypted	Specifies that the password be encrypted using a two-way algorithm.
password	Authentication password you assign.
level { **1** | **2** }	(Optional) Specifies the password for Level 1 or Level 2 independently.

Command Default

Both Level 1 and Level 2 are configured if no level is specified.

Command Modes

Router 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

The **lsp-password accept** command adds an additional password for use when the system validates incoming LSPs and sequence number PDUs (SNPs). An LSP password must be configured using the **lsp-password** command before an accept password can be configured for the corresponding level.

Task ID

Task ID	Operations
isis	read, write

Examples

The following example shows how to configure an Level 1 LSP and SNP password:

```
RP/0/RSP0/CPU0:router(config)# router isis isp
RP/0/RSP0/CPU0:router(config-isis)# lsp-password accept encrypted password1 level 1
```

Related Commands

Command	Description
lsp-password, on page 882	Configures an authentication LSP password.

lsp-refresh-interval

To set the time between regeneration of link-state packets (LSPs) that contain different sequence numbers, use the **lsp-refresh-interval** command in router configuration mode. To restore the default refresh interval, use the **no** form of this command.

lsp-refresh-interval *seconds* [**level** {**1** | **2**}]
no lsp-refresh-interval [*seconds* [**level** {**1** | **2**}]]

Syntax Description

seconds	Refresh interval (in seconds). Range is 1 to 65535 seconds.
level { **1** | **2** }	(Optional) Specifies routing Level 1 or Level 2 independently.

Command Default

seconds : 900 seconds (15 minutes)

Both Level 1 and Level 2 are configured if no level is specified.

Command Modes

Router 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

The refresh interval determines the rate at which the software periodically sends the route topology information that it originates. This behavior is done to keep the information from becoming too old. By default, the refresh interval is 900 seconds (15 minutes).

LSPs must be refreshed periodically before their lifetimes expire. The refresh interval must be less than the LSP lifetime specified with this router command. Reducing the refresh interval reduces the amount of time that undetected link-state database corruption can persist at the cost of increased link utilization. (This event is extremely unlikely, however, because there are other safeguards against corruption.) Increasing the interval reduces the link utilization caused by the flooding of refreshed packets (although this utilization is very small).

Task ID

Task ID	Operations
isis	read, write

Examples

The following example shows how to change the LSP refresh interval to 10,800 seconds (3 hours):

```
RP/0/RSP0/CPU0:router(config)# router isis isp
RP/0/RSP0/CPU0:router(config-isis)# lsp-refresh-interval 10800
```

Related Commands

Command	Description
max-lsp-lifetime, on page 891	Sets the maximum time that LSPs persist without being refreshed.

maximum-paths (IS-IS)

To configure the maximum number of parallel routes that an IP routing protocol will install in the routing table, use the **maximum-paths** command in address family configuration mode. To remove the **maximum-paths** command from the configuration file and restore the system default behavior, use the **no** form of this command. By default up to 8 parallel ECMP paths are used by IS-IS routing protocol.

maximum-paths *maximum*
no maximum-paths

Syntax Description

maximum	Maximum number of parallel routes that IS-IS can install in a routing table. Range is 1 to 64

Command Default

Note The default value used for maximum-paths is 8.

Command Modes

Address family configuration

Command History

Release	Modification
Release 5.3.0	ECMP support extended from 32 to 64 paths.
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 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
isis	read, write

Examples

The following example shows how to allow a maximum of 16 paths to a destination:

```
RP/0/RSP0/CPU0:router(config)# router isis isp
RP/0/RSP0/CPU0:router(config-isis)# address-family ipv4 unicast
RP/0/RSP0/CPU0:router(config-isis-af)# maximum-paths 16
```

maximum-redistributed-prefixes (IS-IS)

To specify an upper limit on the number of redistributed prefixes (subject to summarization) that the Intermediate System-to-Intermediate System (IS-IS) protocol advertises, use the **maximum-redistributed-prefixes** command in address family mode. To disable this feature, use the **no** form of this command.

maximum-redistributed-prefixes *maximum* [**level** {**1** | **2**}]
no maximum-redistributed-prefixes [*maximum* [**level** {**1** | **2**}]]

Syntax Description

maximum	Maximum number of redistributed prefixes advertised. Range is 1 to 28000.
level { **1** | **2** }	(Optional) Specifies maximum prefixes for Level 1 or Level 2.

Command Default

maximum: 10000

level : 1-2

Command Modes

Address family 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **maximum-redistributed-prefixes** command to prevent a misconfiguration from resulting in redistribution of excess prefixes. If IS-IS encounters more than the maximum number of prefixes, it sets a bi-state alarm. If the number of to-be-redistributed prefixes drops back to the maximum or lower—either through reconfiguration or a change in the redistribution source—IS-IS clears the alarm.

Task ID

Task ID	Operations
isis	read, write

Examples

The following example shows how to specify the number of redistributed prefixes at 5000 for Level 2:

```
RP/0/RSP0/CPU0:router(config)# router isis isp
RP/0/RSP0/CPU0:router(config-isis)# address-family ipv4 unicast
RP/0/RSP0/CPU0:router(config-isis-af)# maximum-redistributed-prefixes 5000 level 2
```

maximum-propagated-prefixes (IS-IS)

To configure the maximum number of propagated prefixes, use the **maximum-propagated-prefixes** command in IS-IS configuration mode mode.

maximum-propagated-prefixes *limit* [**warning-only** | **warning-percentage** *percentage*] [**level** *value*]

Syntax Description

maximum-propagated-prefixes *limit*	Configures the maximum number of propagated prefixes.
warning-only	Configures the warning only behavior. The default warning-only threshold is 20,000 prefixes.
warning-percentage *percent*	Configures the warning threshold percentage. The default percentage is 70%
level *value*	Applies the limit to level 1 or level 2 propagation.

Command Default

None

Command Modes

IS-IS configuration

Command History

Release	Modification
Release 26.3.1	This command was introduced.

Usage Guidelines

None

Task ID

Task ID	Operations
isis	read, write

Examples

This example shows how to configure the maximum propagated prefix limit for the level 2 propagation:

```
Router(config)# router isis isp
Router(config-isis)# maximum-propagated-prefixes 8000 level 2
```

This example shows how to set an early warning threshold as a percentage.

```
Router(config)# router isis isp
Router(config-isis)# maximum-propagated-prefixes 8000 warning-percentage 70 level 2
```

This example shows how to configuring a warning-only threshold.

```
Router(config)# router isis isp
Router(config-isis)# maximum-propagated-prefixes 20000 warning-only level 2
```

max-lsp-lifetime

To set the maximum time that link-state packets (LSPs) persist without being refreshed, use the **max-lsp-lifetime** command in router configuration mode. To restore the default time, use the **no** form of this command.

max-lsp-lifetime *seconds* [**level** {**1** | **2**}]
no max-lsp-lifetime [*seconds* [**level** {**1** | **2**}]]

Syntax Description

seconds	Lifetime (in seconds) of the LSP. Range from 1 to 65535 seconds.
level { **1** | **2** }	(Optional) Specifies routing Level 1 or Level 2 independently.

Command Default

seconds : 1200 seconds (20 minutes)

Both Level 1 and Level 2 are configured if no level is specified.

Command Modes

Router 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

You might need to adjust the maximum LSP lifetime if you change the LSP refresh interval with the **lsp-refresh-interval** command. The maximum LSP lifetime must be greater than the LSP refresh interval.

Task ID

Task ID	Operations
isis	read, write

Examples

The following example shows how to set the maximum time that the LSP persists to 11,000 seconds (more than 3 hours):

```
RP/0/RSP0/CPU0:router(config)# router isis isp
RP/0/RSP0/CPU0:router(config-isis)# max-lsp-lifetime 11000
```

Related Commands

Command	Description
lsp-refresh-interval, on page 886	Sets the LSP refresh interval.

max-lsp-tx

To set the maximum number of Link State Packets (LSPs) that can be transmitted per second in an IS-IS routing process, use the **max-lsp-tx** command in IS-IS router lsp-fast-flooding configuration mode. To remove the configured maximum LSP transmission rate, use the **no** form of this command.

max-lsp-tx *lsps-per-sec*

Syntax Description

lsps-per-sec	The maximum number of LSPs that can be transmitted per second. The default is 1000 ms. The range is from 33 through 5000 ms.

Command Default

The command is disabled by default.

Command Modes

IS-IS router lsp-fast-flooding configuration

Command History

Release	Modification
Release 24.3.1	This command was introduced.

Usage Guidelines

None

Task ID

Task ID	Operation
isis	read, write

Example

This example shows how to set maximum LSP transmission rate to 600 LSPs per second.

```
Router# configure terminal
Router(config)# router isis 1
Router(config-isis)# lsp-fast-flooding
Router(config-isis-lsp-fast-flooding)# max-lsp-tx 600
```

max-metric

To configure the Intermediate System-to-Intermediate System (IS-IS) protocol to set maximum metrics for Network Layer Reachability Information (NLRI) during router overload, use the **max-metric** command in IS-IS configuration mode.

max-metric [**on-startup** *timeout* | *wait-for-bgp*] [**default-route** | **external** | **interlevel** | **srv6-locator** [**level**]

Syntax Description

on-startup	Override the default route metric value with the maximum metric value.
timeout	(Optional) Time in seconds to advertise the maximum metric after a reboot.
wait-for-bgp	(Optional) Sets the maximum metric on startup until BGP signal convergence or timeout.
default-route	Overrides the default route metric values with maximum metric value.
external	Overrides the metric of prefixes learned from another protocol with the maximum metric value.
interlevel	Overrides the metric of prefixes learned from another ISIS level with the maximum metric value.
srv6-locator	Overrides the Segment Routing IPv6 metric value with maximum metric value.
level	Sets the maximum metric value for one level only.

Command Default

Maximum metric is disabled.

Command Modes

IS-IS configuration

Command History

Release	Modification
7.0.2	This command was introduced.

Usage Guidelines

The router ignores the **max-metric on-startup** command and keyword combination when the IS-IS process restarts after configuring the **nsf** command, either with **cisco** or **ietf** keyword. This makes the IS-IS Max Metric feature compatible with the hitless switchover and process restart.

Task ID

Task ID	Operations
isis	read, write

Examples

The following example shows how to enable maximum metric on a router:

```
Router(config)# router isis 1
Router(config-isis)# max-metric on-startup wait-for-bgp default-route external interlevel
```

```
srv6-locator level 2
Router(config-isis)# exit
```

metric fallback anomaly

To configure the metric when a delay or loss anomaly occurs for an Intermediate System-to-Intermediate System (IS-IS) interface, use the **metric fallback anomaly** command in address family or interface address family configuration mode.

metric fallback anomaly { [**delay** | **loss**] [**te-metric**] { **increment** *penalty_value* | **multiplier** *penalty_multiplier* | **maximum** } } [**level** { **1** | **2** }]

Syntax Description

maximum	Specifies maximum wide metric. All routers exclude this link from their shortest path first (SPF).
delay	Increases the link metric only when a delay anomaly occurs.
loss	Increases the link metric only when a loss anomaly occurs.
te-metric	Modifies the TE metric.
increment *penalty_value*	Increments the link metric by the specified amount. The range is from 1 to 16777213.
multiplier *penalty_multiplier*	Multiplies the link metric by the specified amount. The range is from 2 to 100000.
level { **1** | **2** }	(Optional) Specifies the SPF calculation for Level 1 or Level 2 independently.

Command Default

Level 1 and Level 2 are configured if no level is specified.

Command Modes

Address family configuration

Interface address family configuration

Command History

Release	Modification
Release 25.1.1	The **loss** keyword was introduced.
Release 7.4.1	This command was introduced.

Usage Guidelines

Specifying the **level** keyword resets the metric only for the specified level. We highly recommend that you configure metrics on all interfaces.

Use the **fallback anomaly** options to increase the metric when you configure link delay anomaly detection in Segment Routing Performance Measurement (SR-PM).

When you configure link delay anomaly detection, SR-PM sends an anomaly bit (A-bit) when the upper bound of the anomaly threshold is exceeded. When an anomaly is received for an interface, IS-IS can increase the IGP metric or TE metric to make the link less attractive for traffic.

From Release 25.1.1, you can configure link loss anomaly detection so that SR-PM sends an anomaly bit (A-bit) when the upper bound of the anomaly threshold is exceeded. When an anomaly is received for an interface, IS-IS can increase the IGP metric or TE metric to make the link less attractive for traffic.

You can specify how to increase the metric in the following ways:

- As a factor by which the base metric is multiplied
- As an increment that is added to the base metric
- As the maximum value. All routers will exclude this link from their SPF.

Task ID

Task ID	Operations
isis	read, write

Examples

In the following example, the base IGP metric value is 10. This example shows how to increase the IGP metrics for an interface. The IGP metric will be increased by 50. When an anomalous bit notification comes for interface HundredGigE 0/1/0/0, the value advertised in the TLVs will be IGP metric 60.

```
Router(config)# router isis 1
Router(config-isis)# interface HundredGigE 0/1/0/0
Router(config-isis-if)# address-family ipv4 unicast
Router(config-isis-if-af)# metric fallback anomaly delay increment 50
```

From Release 25.1.1, you can configure link loss anomaly detection. This example displays how to increase the IGP metrics for an interface.

Consider that the base IGP metric value is 10. You can configure the IGP metric to increase by 500, so that if an anomalous bit notification occurs for interface HundredGigE 0/1/0/0 due to packet loss, the value advertised in the TLVs will be IGP metric 510.

```
Router(config)# router isis 1
Router(config-isis)# interface HundredGigE 0/1/0/0
Router(config-isis-if)# address-family ipv4 unicast
Router(config-isis-if-af)# metric fallback anomaly loss increment 500
```

mesh-group (IS-IS)

To optimize link-state packet (LSP) flooding in highly meshed networks, use the **mesh-group** command in interface configuration mode. To remove a subinterface from a mesh group, use the **no** form of this command.

mesh-group {*number* | **blocked**}
no mesh-group

Syntax Description

number	Number identifying the mesh group of which this interface is a member. Range is 1 to 4294967295.
blocked	Specifies that no LSP flooding takes place on this interface.

Command Default

There is no mesh group configuration (normal LSP flooding).

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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

LSPs first received on subinterfaces that are not part of a mesh group are flooded to all other subinterfaces in the usual way.

LSPs first received on subinterfaces that are part of a mesh group are flooded to all interfaces except those in the same mesh group. If the **blocked** keyword is configured on a subinterface, then a newly received LSP is not flooded out over that interface.

To minimize the possibility of incomplete flooding, you should allow unrestricted flooding over at least a minimal set of links in the mesh. Selecting the smallest set of logical links that covers all physical paths results in very low flooding, but less robustness. Ideally you should select only enough links to ensure that LSP flooding is not detrimental to scaling performance, but enough links to ensure that under most failure scenarios, no router is logically disconnected from the rest of the network. In other words, blocking flooding on all links permits the best scaling performance, but there is no flooding. Permitting flooding on all links results in very poor scaling performance.



Note See RFC 2973 for details about the mesh group specification.

Task ID

Task ID	Operations
isis	read, write

Examples

In the following example, six interfaces are configured in three mesh groups. LSPs received are handled as follows:

- LSPs first received by GigabitEthernet interface 0/1/0/0 are flooded to all interfaces except GigabitEthernet 0/1/0/1 (which is part of the same mesh group) and GigabitEthernet 0/3/0/0 (which is blocked).
- LSPs first received by GigabitEthernet 0/2/0/1 are flooded to all interfaces except GigabitEthernet 0/2/0/0 (which is part of the same mesh group) and GigabitEthernet 0/3/0/0 (which is blocked).
- LSPs first received by GigabitEthernet 0/3/0/0 are not ignored, but flooded as usual to all interfaces.
- LSPs received first through GigabitEthernet 0/3/0/1 are flooded to all interfaces, except GigabitEthernet 0/3/0/0 (which is blocked).

```
RP/0/RSP0/CPU0:router(config)# router isis isp
RP/0/RSP0/CPU0:router(config-isis)# interface GigabitEthernet 0/1/0/0
RP/0/RSP0/CPU0:router(config-isis-if)# mesh-group 10
RP/0/RSP0/CPU0:router(config-isis-if)# exit
RP/0/RSP0/CPU0:router(config-isis)# interface GigabitEthernet 0/1/0/1
RP/0/RSP0/CPU0:router(config-isis-if)# mesh-group 10
RP/0/RSP0/CPU0:router(config-isis-if)# exit
RP/0/RSP0/CPU0:router(config-isis)# interface GigabitEthernet 0/2/0/0
RP/0/RSP0/CPU0:router(config-isis-if)# mesh-group 11
RP/0/RSP0/CPU0:router(config-isis-if)# exit
RP/0/RSP0/CPU0:router(config-isis)# interface GigabitEthernet 0/2/0/1
RP/0/RSP0/CPU0:router(config-isis-if)# mesh-group 11
RP/0/RSP0/CPU0:router(config-isis-if)# exit
RP/0/RSP0/CPU0:routerconfig-isis)# interface GigabitEthernet 0/3/0/1
RP/0/RSP0/CPU0:router(config-isis-if)# mesh-group 12
RP/0/RSP0/CPU0:router(config-isis-if)# exit
RP/0/RSP0/CPU0:router(config-isis)# interface GigabitEthernet 0/3/0/0
RP/0/RSP0/CPU0:router(config-isis-if)# mesh-group blocked
```

metric (IS-IS)

To configure the metric for an Intermediate System-to-Intermediate System (IS-IS) interface, use the **metric** command in address family or interface address family configuration mode. To restore the default metric value, use the **no** form of this command.

metric { *default-metric* | **maximum** | **fallback anomaly** [**delay**] [**te-metric**] { **increment** *penalty_value* | **multiplier** *penalty_multiplier* | **maximum** } } [**level** { **1** | **2** }]

Syntax Description

default-metric	Metric assigned to the link and used to calculate the cost from each other router using the links in the network to other destinations. Range is 1 to 63 for narrow metric and 1 to 16777214 for wide metric. **Note** Setting the default metric under address family results in setting the same metric for all interfaces that is associated with the address family. Setting a metric value under an interface overrides the default metric
maximum	Specifies maximum wide metric. All routers exclude this link from their shortest path first (SPF).
fallback	Increases the link metric when an event occurs.
anomaly	Increases the link metric when a delay or loss anomaly occurs.
delay	Increases the link metric only when a delay anomaly occurs.
te-metric	Modifies the TE metric.
increment <penalty_value>	Increments the link metric by the specified amount. The range is from 1 to 16777213.
multiplier <penalty_multiplier>	Multiplies the link metric by the specified amount. The range is from 2 to 100000.
level { **1** | **2** }	(Optional) Specifies the SPF calculation for Level 1 or Level 2 independently.

Command Default

default-metric : Default is 10.

Both Level 1 and Level 2 are configured if no level is specified.

Command Modes

Address family configuration

Interface address family configuration

Command History

Release	Modification
Release 7.4.1	The **fallback anomaly** keywords was introduced.
Release 3.7.2	This command was introduced.

Usage Guidelines

Specifying the **level** keyword resets the metric only for the specified level. We highly recommend that you configure metrics on all interfaces.

Set the default metric under address family to set the same metric for all interfaces that is associated with the address family. Set a metric value under an interface to override the default metric.

We highly recommend that you configure metrics on all interfaces.

Metrics of more than 63 cannot be used with narrow metric style.

Use the **fallback anomaly** options to increase the metric when you configure link delay anomaly detection in Segment Routing Performance Measurement (SR-PM).

When you configure link delay anomaly detection, SR-PM sends an anomaly bit (A-bit) when the upper bound of the anomaly threshold is exceeded. When an anomaly is received for an interface, IS-IS can increase the IGP metric or TE metric to make the link less attractive for traffic.

You can specify how to increase the metric in the following ways:

- As a factor by which the base metric is multiplied
- As an increment that is added to the base metric
- As the maximum value. All routers will exclude this link from their SPF.

Task ID

Task ID	Operations
isis	read, write

Examples

The following example shows how to configure Packet-over-SONET/SDH 0/1/0/1 interface with a default link-state metric cost of 15 for Level 1:

```
Router(config)# router isis isp
Router(config-isis)# interface GigabitEthernet /1/0/1
Router(config-isis-if)# address-family ipv4 unicast
Router(config-isis-if-af)# metric 15 level 1
```

The following example shows how to configure a metric cost of 15 for all interfaces under address family IPv4 unicast for level 2:

```
Router# configure
Router(config)# router isis isp
Router(config-isis)# address-family ipv4 unicast
Router(config-isis-af)# metric 15 level 2
```

In the following example, the base IGP metric value is 10. This example shows how to increase the IGP metrics for an interface. The IGP metric will be increased by 50. When an anomalous bit notification comes for interface HundredGigE 0/1/0/0, the value advertised in the TLVs will be IGP metric 60.

```
Router(config)# router isis 1
Router(config-isis)# interface HundredGigE 0/1/0/0
```

```
Router(config-isis-if)# address-family ipv4 unicast
Router(config-isis-if-af)# metric fallback anomaly delay increment 50
```

Related Commands

Command	Description
metric-style narrow, on page 902	Configures a router running IS-IS so that it generates and accepts old-style TLV objects.
metric-style transition, on page 904	Configures the software to generate and accept both old-style and new-style TLV objects.
metric-style wide, on page 906	Configures the software to generate and accept only new-style TLV objects objects.

metric-style narrow

To configure the Intermediate System-to-Intermediate System (IS-IS) software to generate and accept old-style type, length, and value (TLV) objects, use the **metric-style narrow** command in address family configuration mode. To remove the **metric-style narrow** command from the configuration file and restore the system to its default condition, use the **no** form of this command.

metric-style narrow [**transition**] [**level** {**1** | **2**}]
no metric-style narrow [**transition**] [**level** {**1** | **2**}]

Syntax Description

transition	(Optional) Instructs the router to generate and accept both old-style and new-style TLV objects. It generates only old-style TLV objects.
level { **1** | **2** }	(Optional) Specifies routing Level 1 or Level 2 independently.

Command Default

Old-style TLVs are generated.

Both Level 1 and Level 2 are configured if no level is specified.

Command Modes

Address family 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

IS-IS traffic engineering extensions include new-style TLV objects with wider metric fields than old-style TLV objects. By default, the router generates old-style TLV objects only. To perform Multiprotocol Label Switching traffic engineering (MPLS TE), a router must generate new-style TLV objects.

Task ID

Task ID	Operations
isis	read, write

Examples

The following example shows how to configure the router to generate and accept only old-style TLV objects on router Level 1:

```
RP/0/RSP0/CPU0:router(config)# router isis isp
RP/0/RSP0/CPU0:router(config-isis)# address-family ipv4 unicast
RP/0/RSP0/CPU0:router(config-isis-af)# metric-style narrow level 1
```

Related Commands

Command	Description
metric-style transition, on page 904	Configures a router to generate and accept both old-style and new-style TLV objects.
metric-style wide, on page 906	Configures a router to generate and accept only new-style TLV objects.

metric-style transition

To configure the Intermediate System-to-Intermediate System (IS-IS) software to generate and accept both old-style and new-style type, length, and value (TLV) objects, use the **metric-style transition** command in address family configuration mode. To remove the **metric-style transition** command from the configuration file and restore the system to its default condition, use the **no** form of this command.

metric-style transition [**level** {**1** | **2**}]
no metric-style transition [**level** {**1** | **2**}]

Syntax Description

transition	Instructs the router to generate and accept both old-style and new-style TLV objects.
level { **1** | **2** }	(Optional) Specifies routing Level 1 or Level 2 independently.

Command Default

Old-style TLVs are generated, if this command is not configured.

Both Level 1 and Level 2 are configured if no level is specified.

Command Modes

Address family 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

IS-IS traffic engineering extensions include new-style TLV objects which have wider metric fields than old-style TLV objects. By default, the router generates old-style TLV objects only. To perform Multiprotocol Label Switching traffic engineering (MPLS TE), a router needs to generate new-style TLV objects.

Task ID

Task ID	Operations
isis	read, write

Examples

The following example shows how to configure the router to generate and accept both old-style and new-style TLV objects on Level 2:

```
RP/0/RSP0/CPU0:router(config)# router isis isp
RP/0/RSP0/CPU0:router(config-isis)# address-family ipv4 unicast
RP/0/RSP0/CPU0:router(config-isis-af)# metric-style transition level 2
```

Related Commands

Command	Description
metric-style narrow, on page 902	Configures a router to generate and accept only old-style TLV objects.
metric-style wide, on page 906	Configures a router to generate and accept only new-style TLV objects.

metric-style wide

To configure the Intermediate System-to-Intermediate System (IS-IS) software to generate and accept only new-style type, length, and value (TLV) objects, use the **metric-style wide** command in address family configuration mode. To remove the **metric-style wide** command from the configuration file and restore the system to its default condition, use the **no** form of this command.

metric-style wide [**transition**] [**level** {**1** | **2**}]
no metric-style wide [**transition**] [**level** {**1** | **2**}]

Syntax Description

transition	(Optional) Instructs the router to generate and accept both old-style and new-style TLV objects. It generates only new-style TLV objects.
level { **1** | **2** }	(Optional) Specifies routing Level 1 or Level 2 independently.

Command Default

Old-style TLV lengths are generated, if this command is not configured.

Both Level 1 and Level 2 are configured if no level is specified.

Command Modes

Address family 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

IS-IS traffic engineering extensions include new-style TLV objects with wider metric fields than old-style TLV objects. If you enter the **metric-style wide** command, a router generates and accepts only new-style TLV objects. Therefore, the router uses less memory and fewer other resources rather than generating both old-style and new-style TLV objects.

To perform MPLS traffic engineering, a router needs to generate new-style TLV objects.

Note This discussion of metric styles and transition strategies is oriented toward traffic engineering deployment. Other commands and models might be appropriate if the new-style TLV objects are desired for other reasons. For example, a network may require wider metrics, but might not use traffic engineering.

Task ID

Task ID	Operations
isis	read, write

Examples

The following example shows how to configure a router to generate and accept only new-style TLV objects on Level 1:

```
RP/0/RSP0/CPU0:router(config)# router isis isp
RP/0/RSP0/CPU0:router(config-isis)# address-family ipv4 unicast
RP/0/RSP0/CPU0:router(config-isis-af)# metric-style wide level 1
```

Related Commands

Command	Description
metric-style narrow, on page 902	Configures a router to generate and accept only old-style TLV objects.

microloop avoidance (IS-IS)

To avoid micro-loops by delaying the convergence of all or protected prefixes, use the **microloop avoidance** command. Valid triggers for microloop avoidance feature are local link-down events only, such as link down, BFD down, and IS-IS adjacency down. Microloops caused by other triggers are not avoided by this feature. Consider `microloop avoidance segment-routing` command for extended trigger coverage.

To disable this function, use the **no** prefix for this command.

microloop avoidance [**protected** | **rib-update-delay** *delay*]
no microloop avoidance

Syntax Description

(none)	Delays convergence of all prefixes.
protected	(Optional) Delays convergence of protected prefixes.
rib-update-delay *delay*	(Optional) Delays convergence of all prefixes and updates RIB after the configured delay. The range is 1 to 60000 milliseconds. The default is 5000ms (for both the flavours of uloop avoidance).

Command Default

Micro-loop avoidance is disabled by default.

Command Modes

router isis configuration

Command History

Release	Modification
Release 4.3.1	This command was introduced.

Usage Guidelines

When the network converges after a link failure restoration, micro-loops can form due to inconsistencies in the forwarding tables of different routers. By delaying the convergence of prefixes, you can avoid the formation of micro-loops.

You can delay the convergence of all or protected prefixes by using the **microloop avoidance** command. When configured, the command applies to all prefixes by default. To enable it for only protected prefixes, use the **protected** option.

If another event occurs when the microloop avoidance timer is running, the microloop avoidance process is cancelled, and RIB delay timer is cancelled and prefixes are sent to RIB immediately.

Task ID

Task ID	Operations
isis	read, write

Examples

The following example shows how to configure micro-loop avoidance with IS-IS:

```
Router# configure
```

```
Router(config)# router isis 50
Router(config-isis)# microloop avoidance rib-update-delay 400
```

min-lsp-arrivaltime

To control the rate of incoming LSPs (link-state packets) LSPs, use the **min-lsp-arrivaltime** command in router configuration mode. To remove this function use the **no** form of this command.

min-lsp-arrivaltime [**initial-wait initial**] [**secondary-wait secondary**] [**maximum-wait maximum**] [**level** {**1** | **2**}]
no min-lsp-arrivaltime [**initial-wait initial**] [**secondary-wait secondary**] [**maximum-wait maximum**] [**level** {**1** | **2**}]

Syntax Description

initial-wait initial	Initial LSP calculation delay (in milliseconds). Range is 0 to 120000.
secondary-wait *secondary*	Hold time between the first and second LSP calculations (in milliseconds). Range is 0 to 120000.
maximum-wait *maximum*	Maximum interval (in milliseconds) between two consecutive LSP calculations. Range is 0 to 120000.
level { **1** | **2** }	(Optional) Enables the LSP interval configuration for Level 1 or Level 2 independently.

Command Default

Both Level 1 and Level 2 are configured if no level is specified.

Command Modes

Router configuration mode

Command History

Release	Modification
Release 3.9.0	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.

This command can be used to protect a router against the possible instability of its neighbor's LSPs.

The command parameters are similair to **lsp-gen-interval**command and neighbors lsp-gen-interval values can be used to set the **min-lsp-arrivaltime**

Note The initial-wait of minimum-lsp-arrival has no use in computing maximum counts and maximum window sizes of the LSP arrival time parameter.

Task ID

Task ID	Operations
isis	read, write

Examples

The following example shows how to configure min-lsp-arrival time commands:

```
RP/0/RSP0/CPU0:router(config)# router isis isp
RP/0/RSP0/CPU0:router(config)# router isis isp min-lsp-arrivaltime
RP/0/RSP0/CPU0:router(config)# router isis 1 min-  lsp-arrivaltime initial-wait
RP/0/RSP0/CPU0:router(config)#router isis 1 min-lsp-arrivaltime maximum-wait
RP/0/RSP0/CPU0:router(config)#router isis 1 min-lsp-arrivaltime secondary-wait
```

mpls ldp auto-config

To enable Label Distribution Protocol (LDP) Interior Gateway Protocol (IGP) interface auto-configuration, use the **mpls ldp auto-config** command in IPv4 address family configuration mode. To disable LDP IGP auto-configuration, use the **no** form of this command.

mpls ldp auto-config
no mpls ldp auto-config

Syntax Description

This command has no keywords or arguments.

Command Default

LDP IGP auto-configuration is disabled.

Command Modes

IPv4 address family 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **mpls ldp auto-config** command to automatically configure LDP on a set of interfaces associated with a specified IGP instance. Further, LDP IGP auto-configuration provides a means to block LDP from being enabled on a specified interface. If you do not want an IS-IS interface to have LDP enabled, use the **igp auto-config disable** command.

Task ID

Task ID	Operations
isis	read, write

Examples

The following example shows how to enable LDP IGP auto-configuration:

```
RP/0/RSP0/CPU0:router(config)# router isis isp
RP/0/RSP0/CPU0:router(config-isis)# address-family ipv4 unicast
RP/0/RSP0/CPU0:router(config-isis-af)# mpls ldp auto-config
```

Related Commands

Command	Description
igp auto-config disable	Disables LDP IGP auto-configuration for a specific interface.

mpls ldp sync (IS-IS)

To configure Label Distribution Protocol (LDP) IS-IS synchronization, use the **mpls ldp sync** command in interface address family configuration mode. To disable LDP synchronization, use the **no** form of this command.

mpls ldp sync [**level** {**1** | **2**}]
no mpls ldp sync [**level** {**1** | **2**}]

Syntax Description

level { **1 | 2** }	(Optional) Sets LDP synchronization for the specified level.

Command Default

If a level is not specified, LDP synchronization is set for both levels.

Command Modes

Interface address family 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

MPLS VPN traffic forwarded using LDP labels can be dropped in the following instances:

- A new link is introduced in the network and IS-IS has converged before LDP establishes labels.
- An existing LDP session goes down while IS-IS adjacency is intact over the link.

In both instances, outbound LDP labels are not available for forwarding MPLS traffic. LDP IS-IS synchronization addresses the traffic drop. When the **mpls ldp sync** command is configured, IS-IS advertises the maximum possible link metric until LDP has converged over the link. The link is less preferred and least used in forwarding MPLS traffic. When LDP establishes the session and exchanges labels, IS-IS advertises the regular metric over the link.



Note IS-IS advertises the maximum metric –1 (16777214) if wide metrics are configured since the maximum wide metric is specifically used for link exclusion from the shortest path first algorithm (SPF) (RFC 3784). However, the maximum narrow metric is unaffected by this definition.

Task ID

Task ID	Operations
isis	read, write

Examples

The following example shows how to enable LDP IS-IS synchronization:

```
RP/0/RSP0/CPU0:router(config)# router isis isp
RP/0/RSP0/CPU0:router(config-isis)# interface GigabitEthernet 0/3/0/0
RP/0/RSP0/CPU0:router(config-isis-if)# address-family ipv4 unicast
RP/0/RSP0/CPU0:router(config-isis-if-af)# mpls ldp sync
```

Related Commands

Command	Description
show isis interface, on page 979	Displays information about the IS-IS interfaces

mpls traffic-eng (IS-IS)

To configure a router running the Intermediate System-to-Intermediate System (IS-IS) protocol to flood Multiprotocol Label Switching traffic engineering (MPLS TE) link information into the indicated IS-IS level, use the **mpls traffic-eng** command in IPv4 address family configuration mode. To disable this feature, use the **no** form of this command.

mpls traffic-eng {**level-1** | **level-1-2** | **level-2-only**}
no mpls traffic-eng [**level-1** | **level-1-2** | **level-2-only**]

Syntax Description

level-1	Specifies routing level 1.
level-1-2	Specifies routing levels 1 and 2.
level-2-only	Specifies routing level 2.

Command Default

Flooding is disabled.

Command Modes

IPv4 address family 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **mpls traffic-eng** command, which is part of the routing protocol tree, to flood link resource information (such as available bandwidth) for appropriately configured links in the link-state packet (LSP) of the router.

Task ID

Task ID	Operations
isis	read, write

Examples

The following example shows how to turn on MPLS traffic engineering for IS-IS level 1:

```
RP/0/RSP0/CPU0:router(config)# router isis isp
RP/0/RSP0/CPU0:router(config-isis)# address-family ipv4 unicast
RP/0/RSP0/CPU0:router(config-isis-af)# mpls  traffic-eng level-1
```

Related Commands

Command	Description
mpls traffic-eng router-id (IS-IS), on page 920	Specifies that the traffic engineering router identifier for the node is the IP address associated with a given interface.

mpls traffic-eng multicast-intact (IS-IS)

To enable multicast-intact for Intermediate System-to-Intermediate System (IS-IS) routes with Protocol-Independent Multicast (PIM) and Multiprotocol Label Switching (MPLS) traffic engineering, use the **mpls traffic-eng multicast-intact** command in IPv4 address family configuration mode. To disable this feature, use the **no** form of this command.

mpls traffic-eng multicast-intact
no mpls traffic-eng [**multicast-intact**]

Syntax Description

This command has no keywords or arguments.

Command Default

Multicast-intact is disabled.

Command Modes

IPv4 address family 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

If Multiprotocol Label Switching Traffic Engineering (MPLS-TE) is configured through the IS-IS routing domain and multicast protocols (like Protocol Independent Multicast [PIM]) are also enabled, then use the **mpls traffic-end multicast-intact** command to install nontraffic engineering next hops in the Routing Information Base (RIB) for use by multicast. The installation of IP-only next hops is in addition to the installation of the standard set of paths for a prefix, which might be through traffic engineered tunnels.

The **mpls traffic-eng multicast-intact** command allows PIM to use the native hop-by-hop neighbors even though the unicast routing is using MPLS TE tunnels.

Examples

The following example shows how to enable the multicast-intact feature:

```
RP/0/RSP0/CPU0:router(config)# router isis isp
RP/0/RSP0/CPU0:router(config-isis)# address-family ipv4 unicast
RP/0/RSP0/CPU0:router(config-isis-af)# mpls  traffic-engmulticast-intact
```

Related Commands

Command	Description
show isis route, on page 998	Displays IP reachability information for an IS-IS instance, optionally for multicast-intact.
show isis topology, on page 1012	Displays a list of connected IS-IS routers in all areas, optionally for multicast-intact.

mpls traffic-eng path-selection ignore overload

To ensure that label switched paths (LSPs) are not disabled when routers have the Intermediate System-to-Intermediate System (IS-IS) overload bit set, use the **mpls traffic-eng path-selection ignore overload** command in global configuration mode. To disable this override, use the **no** form of this command.

mpls traffic-eng path-selection ignore overload
no mpls traffic-eng path-selection ignore overload

Command Default

No default behavior or values

Command Modes

Global 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

When the IS-IS overload bit avoidance feature is activated, which means that they are still available for use label switched paths (LSPs), all nodes with the overload bit set, including the following nodes, are ignored:

- head nodes
- mid nodes
- tail nodes

Task ID

Task ID	Operations
mpls-te	read, write

Examples

The following example shows how to activate IS-IS overload bit avoidance:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# mpls traffic-eng path-selection ignore overload
```

The following example shows how to deactivate IS-IS overload bit avoidance:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# no mpls traffic-eng path-selection ignore overload
```

Related Commands

Command	Description
set-overload-bit, on page 955	Configures a router to signal other routers not to use it as an intermediate hop in their shortest path first (SPF) calculations.

mpls traffic-eng router-id (IS-IS)

To specify the Multiprotocol Label Switching traffic engineering (MPLS TE) router identifier for the node, use the **mpls traffic-eng router-id** command in IPv4 address family configuration mode. To disable this feature, use the **no** form of this command.

mpls traffic-eng router-id {*ip-address* | *type interface-path-id*}
no mpls traffic-eng [**router-id**]

Syntax Description

ip-address	IP address in four-part, dotted-decimal notation.
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

Global router identifier is used.

Command Modes

IPv4 address family 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

The identifier of the router acts as a stable IP address for the traffic engineering configuration. This IP address is flooded to all nodes. For all traffic engineering tunnels originating at other nodes and ending at this node, you must set the tunnel destination to the traffic engineering router ID of the destination node, because that is the address used by the traffic engineering topology database at the tunnel head for its path calculation.

Note We recommend that loopback interfaces be used for MPLS TE, because they are more stable than physical interfaces.

Task ID

Task ID	Operations
isis	read, write

Examples

The following example shows how to specify the traffic engineering router identifier as the IP address associated with loopback interface 0:

```
RP/0/RSP0/CPU0:router(config)# router isis isp
RP/0/RSP0/CPU0:router(config-isis)# address-family ipv4 unicast
RP/0/RSP0/CPU0:router(config-isis-af)# mpls traffic-eng router-id Loopback0
```

Related Commands

Command	Description
mpls traffic-eng (IS-IS), on page 915	Turns on flooding of MPLS traffic engineering link information in the indicated IGP level or area.

net

To configure an Intermediate System-to-Intermediate System (IS-IS) network entity title (NET) for the routing instance, use the **net** command in router configuration mode. To remove the **net** command from the configuration file and restore the system to its default condition, use the **no** form of this command.

net *network-entity-title*
no net *network-entity-title*

Syntax Description

network-entity-title	NET that specifies the area address and the system ID for an ISIS routing process.

Command Default

No NET is configured. The IS-IS instance is not operational, because a NET is mandatory.

Command Modes

Router configuration

Command History

Release	Modification
Release 3.7.2	This command was introduced.

Usage Guidelines

Under most circumstances, one and only one NET should be configured.

A NET is a network service access point (NSAP) where the last byte is always 0. On a Cisco router running IS-IS, a NET can be 8 to 20 bytes in length. The last byte is always the n-selector and must be 0. The n-selector indicates to which transport entity the packet is sent. An n-selector of 0 indicates no transport entity and means that the packet is for the routing software of the system.

The six bytes directly preceding the n-selector are the system ID. The system ID length is a fixed size and cannot be changed. The system ID must be unique throughout each area (Level 1) and throughout the backbone (Level 2).

All bytes preceding the system ID are the area ID.

A maximum of three NETs for each router is allowed. In rare circumstances, it is possible to configure two or three NETs. In such a case, the area this router is in has three area addresses. Only one area still exists, but it has more area addresses.

Configuring multiple NETs can be temporarily useful in network reconfiguration in which multiple areas are merged, or in which one area is split into more areas. Multiple area addresses enable you to renumber an area individually as needed.

Task ID

Task ID	Operations
isis	read, write

Examples

The following example shows how to configure a router with NET area ID 47.0004.004d.0001 and system ID 0001.0c11.1110:

```
RP/0/RSP0/CPU0:router(config)# router isis isp
RP/0/RSP0/CPU0:router(config-isis)# net 47.0004.004d.0001.0001.0c11.1110.00
```

Related Commands

Command	Description
log adjacency changes (IS-IS), on page 872	Configures the routing level for an instance of the IS-IS routing process.
router isis, on page 950	Enables the IS-IS routing protocol and specifies an IS-IS instance.

nsf (IS-IS)

To enable nonstop forwarding (NSF) on the next restart, use the **nsf** command in router configuration mode. To restore the default setting, use the **no** form of this command.

nsf {**cisco** | **ietf**}
no nsf {**cisco** | **ietf**}

Syntax Description

cisco	Specifies Cisco-proprietary NSF restart.
ietf	Specifies Internet Engineering Task Force (IETF) NSF restart.

Command Default

NSF is disabled.

Command Modes

Router 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

NSF allows an Intermediate System-to-Intermediate System (IS-IS) instance to restart using checkpointed adjacency and link-state packet (LSP) information, and to perform restart with no impact on its neighbor routers. In other words, there is no impact on other routers in the network due to the destruction and recreation of adjacencies and the system LSP.

Task ID

Task ID	Operations
isis	read, write

Examples

The following example shows how to enable Cisco proprietary NSF:

```
RP/0/RSP0/CPU0:router(config)# router isis isp
RP/0/RSP0/CPU0:router(config-isis)# nsf cisco
```

Related Commands

Command	Description
nsf interface-expires, on page 926	Configures the number of resends of an an acknowledged NSF-restart acknowledgment.

Command	Description
nsf interface-timer, on page 928	Configures the time interval after which an unacknowledged IETF NSF restart attempt is repeated.
nsf lifetime (IS-IS), on page 930	Configures the maximum route lifetime following an NSF restart.

nsf interface-expires

To configure the number of resends of an acknowledged nonstop forwarding (NSF)-restart acknowledgment, use the **nsf interface-expires** command in router configuration mode. To restore the default value, use the **no** form of this command.

nsf interface-expires *number*
no nsf interface-expires

Syntax Description

number	Number of resends. Range is 1 to 3.

Command Default

number : 3 resends

Command Modes

Router 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

When a hello packet sent with the NSF restart flag set is not acknowledged, it is re-sent. Use the **nsf interface-expires** command to control the number of times the NSF hello is re-sent. When this limit is reached on an interface, any neighbor previously known on that interface is assumed to be down and the initial shortest path first (SPF) calculation is permitted, provided that all other necessary conditions are met.

The total time period available for adjacency reestablishment (interface-timer * interface-expires) should be greater than the expected total NSF restart time.

The **nsf interface-expires** command applies only to Internet Engineering Task Force (IETF)-style NSF. It has no effect if Cisco-proprietary NSF is configured.

Task ID

Task ID	Operations
isis	read, write

Examples

The following example shows how to allow only one retry attempt on each interface if an IETF NSF restart signal is not acknowledged:

```
RP/0/RSP0/CPU0:router(config)# router isis isp
RP/0/RSP0/CPU0:router(config-isis)# nsf ietf
RP/0/RSP0/CPU0:router(config-isis)# nsf interface-expires 1
```

Related Commands

Command	Description
hello-multiplier, on page 848	Specifies the number of IS-IS hello packets a neighbor must miss before the router should declare the adjacency as down.
nsf interface-timer, on page 928	Configures the time interval after which an unacknowledged IETF NSF restart attempt is repeated.

nsf interface-timer

To configure the time interval after which an unacknowledged Internet Engineering Task Force (IETF) nonstop forwarding (NSF) restart attempt is repeated, use the **nsf interface-timer** command in router configuration mode. To restore the default value, use the **no** form of this command.

nsf interface-timer *seconds*
no nsf interface-timer

Syntax Description

seconds	NSF restart time interval (in seconds). Range is 3 to 20 seconds.

Command Default

seconds : 10 seconds

Command Modes

Router 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

When the IETF NSF restart process begins, hello packets send an NSF restart flag that must be acknowledged by the neighbors of the router. Use the **nsf interface-timer** command to control the restart time interval after the hello packet is re-sent. The restart time interval need not match the hello interval.

The **nsf interface-timer** command applies only to IETF-style NSF. It has no effect if Cisco proprietary NSF is configured.

Task ID

Task ID	Operations
isis	read, write

Examples

The following example shows how to ensure that a hello packet with the NSF restart flag set is sent again every 5 seconds until the flag is acknowledged:

```
RP/0/RSP0/CPU0:router(config)# router isis isp
RP/0/RSP0/CPU0:router(config-isis)# nsf ietf
RP/0/RSP0/CPU0:router(config-isis)# nsf interface-timer 5
```

Related Commands

Command	Description
nsf interface-expires, on page 926	Configures the number of resends of an acknowledged NSF-restart acknowledgment.
hello-interval (IS-IS), on page 846	Specifies the length of time between hello packets that the software sends.

nsf lifetime (IS-IS)

To configure the maximum route lifetime following a nonstop forwarding (NSF) restart, use the **nsf lifetime** command in router configuration mode. To restore the default value, use the **no** form of this command.

nsf lifetime *seconds*
no nsf lifetime

Syntax Description

seconds	Maximum route lifetime (in seconds) following an NSF restart. Range is 5 to 300 seconds.

Command Default

seconds : 60 seconds (1 minute)

Command Modes

Router 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **nsf lifetime** command to set the maximum available time for the reacquisition of checkpointed adjacencies and link-state packets (LSPs) during a Cisco proprietary NSF restart. LSPs and adjacencies not recovered during this time period are abandoned, thus causing changes to the network topology.

Task ID

Task ID	Operations
isis	read, write

Examples

The following example shows how to configure the router to allow only 20 seconds for the entire NSF process:

```
RP/0/RSP0/CPU0:router(config)# router isis isp
RP/0/RSP0/CPU0:router(config-isis)# nsf cisco
RP/0/RSP0/CPU0:router(config-isis)# nsf lifetime 20
```

oor-set-overload-bit disable

To prevent the setting of the overload bit during out-of-resource (OOR) conditions, use **oor-set-overload-bit disable** command in router configuration mode. To allow the setting of the overload bit during OOR conditions, use the **no** form of the command.

oor-set-overload-bit disable

no oor-set-overload-bit disable

Syntax Description

This command has no keywords or arguments.

Command Default

The overload bit is set.

Command Modes

Router configuration

Command History

Release	Modification
Release 24.2.1	This command was introduced.

Usage Guidelines

The setting of the overload bit is typically allowed only when a router experiences issues, such as memory shortages, which might result in an incomplete link-state database and inaccurate routing tables. When the overload bit is set in the link-state packets (LSPs) of an unreliable router, other routers can ignore this router in their shortest path first (SPF) calculations until it recovers. Consequently, no paths through the unreliable router are visible to other routers in the IS-IS area, though IP prefixes directly connected to this router remain reachable.

We recommend not disabling the setting of the overload-bit during OOR conditions. IS-IS will continue to forward traffic but will base decisions on an incomplete link-state database, potentially causing routing loops and other forwarding failures.

Task ID

Task ID	Operation
isis	read, write

Example

The following example shows how to disable the overload bit:

```
Router# router isis 1
Router# (config-isis)# oor-set-overload-bit disable
```

passive (IS-IS)

To suppress Intermediate System-to-Intermediate System (IS-IS) packets from being transmitted to the interface and received packets from being processed on the interface, use the **passive** command in interface configuration mode. To restore IS-IS packets coming to an interface, use the **no** form of this command.

passive
no passive

Command Default

Interface is active.

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 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
isis	read, write

Examples

The following example shows how to configure the router to suppress IS-IS packets on GigabitEthernet interface 0/1/0/1:

```
RP/0/RSP0/CPU0:router(config)# router isis isp
RP/0/RSP0/CPU0:router(config-isis)# interface GigabitEthernet 0/1/0/1
RP/0/RSP0/CPU0:router(config-isis-if)# passive
```

Related Commands

Command	Description
suppressed, on page 1030	Allows the IS-IS interface to participate in forming adjacencies without advertising connected prefixes in the LSPs.

point-to-point

To configure a network of only two networking devices that use broadcast media and the integrated Intermediate System-to-Intermediate System (IS-IS) routing protocol to function as a point-to-point link instead of a broadcast link, use the **point-to-point** command in interface configuration mode. To disable the point-to-point usage, use the **no** form of this command.

point-to-point
no point-to-point

Syntax Description

This command has no keywords or arguments.

Command Default

Interface is treated as broadcast if connected to broadcast media.

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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **point-to-point** command only on broadcast media in a network with two networking devices. The command causes the system to issue packets point-to-point rather than as broadcasts. Configure the command on both networking devices in the network.

Task ID

Task ID	Operations
isis	read, write

Examples

The following example shows how to configure a 10-Gb Ethernet interface to act as a point-to-point interface:

```
RP/0/RSP0/CPU0:router(config)# router isis isp
RP/0/RSP0/CPU0:router(config-isis)# interface TenGigE 0/6/0/0
RP/0/RSP0/CPU0:router(config-isis-if)# point-to-point
```

priority (IS-IS)

To configure the priority of designated routers, use the **priority** command in interface configuration mode. To reset the default priority, use the **no** form of this command.

priority *value* [**level** {**1** | **2**}]
no priority [*value*] [**level** {**1** | **2**}]

Syntax Description

value	Priority of a router. Range is 0 to 127.
level { **1** | **2** }	(Optional) Specifies routing Level 1 or Level 2 independently.

Command Default

value : 64

Both Level 1 and Level 2 are configured if no level is specified.

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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Priorities can be configured for Level 1 and Level 2 independently. Specifying Level 1 or Level 2 resets priority only for Level 1 or Level 2 routing, respectively. Specifying no level allows you to configure all levels.

The priority is used to determine which router on a LAN is the designated router or Designated Intermediate System (DIS). The priorities are advertised in the hello packets. The router with the highest priority becomes the DIS.

In the Intermediate System-to-Intermediate System (IS-IS) protocol, there is no backup designated router. Setting the priority to 0 lowers the chance of this system becoming the DIS, but does not prevent it. If a router with a higher priority comes online, it takes over the role from the current DIS. For equal priorities, the higher MAC address breaks the tie.

Task ID

Task ID	Operations
isis	read, write

Examples

The following example shows how to give Level 1 routing priority by setting the priority level to 80. This router is now more likely to become the DIS.

```
RP/0/RSP0/CPU0:router(config)# router isis isp
```

```
RP/0/RSP0/CPU0:router(config-isis)# interface TenGigE0/6/0/0
RP/0/RSP0/CPU0:router(config-isis-if)# priority 80 level 1
```

propagate level

To propagate routes from one Intermediate System-to-Intermediate System (IS-IS) level into another level, use the **propagate level** command in address family configuration mode. To disable propagation, use the **no** form of this command.

propagate level {**1** | **2**} **into level** {**1** | **2**} **route-policy** *route-policy-name*
no propagate level {**1** | **2**} **into level** {**1** | **2**}

Syntax Description

level { **1** | **2** }	Propagates from routing Level 1 or Level 2 routes.
into	Propagates from Level 1 or Level 2 routes into Level 1 or Level 2 routes.
route-policy *route-policy-name*	Specifies a configured route policy.

Command Default

Route leaking (Level 2 to Level 1) is disabled.

Command Modes

Address family 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

In general, route propagation from Level 1 to Level 2 is automatic. You might want to use this command to better control which Level 1 routes can be propagated into Level 2.

Propagating Level 2 routes into Level 1 is called *route leaking*. Route leaking is disabled by default. That is, Level 2 routes are not automatically included in Level 1 link-state packets (LSPs). If you want to leak Level 2 routes into Level 1, you must enable that behavior by using this command.

Propagation from Level 1 into Level 1 and from Level 2 into Level 2 is not allowed.

Task ID

Task ID	Operations
isis	read, write

Examples

The following example shows how to redistribute Level 2 routes to Level 1:

```
RP/0/RSP0/CPU0:router(config)# ipv4 access-list 101 permit ip 10.0.0.0 255.0.0.0 10.1.0.1
0.255.255.255
RP/0/RSP0/CPU0:router(config)# router isis isp
RP/0/RSP0/CPU0:router(config-isis)# net 49.1234.2222.2222.2222.00
RP/0/RSP0/CPU0:router(config-isis)# address-family ipv4 unicast
RP/0/RSP0/CPU0:router(config-isis-af)# propagate level 2 into level 1 route-policy policy_a
```

Related Commands

Command	Description
redistribute (IS-IS), on page 940	Redistributes routes from one routing domain into a specified IS-IS instance.

protocol shutdown (IS-IS)

To disable an instance of the Intermediate System-to-Intermediate System (IS-IS) protocol, so that it cannot form an adjacency on any interface, use the **protocol shutdown** command in the isis configuration mode. To re-enable the IS-IS protocol, use the **no** form of this command.

protocol shutdown
no protocol shutdown

Command Default

None

Command Modes

isis configuration

Command History

Release	Modification
Release 24.4.1	This command was introduced.

Usage Guidelines

Use the **protocol shutdown** command to disable the IS-IS protocol for a specific routing instance without removing any existing IS-IS configuration parameters.

This command is similar to executing the following sample configuration:

```
router isis 1
 set-overload-bit
 exit
router isis 1
 interface HundredGigE 0/1/0/3
  shutdown
```

Task ID

Task ID	Operations
isis	read, write

Examples

This example shows how to disable the IS-IS.

```
Router# config
Router(config)# router isis 1
Router(config-isis)# protocol shutdown
```

psnp-interval

To configure the interval at which the Intermediate System to Intermediate System (IS-IS) protocol sends Partial Sequence Number PDUs (PSNPs) to acknowledge received Link State PDUs (LSPs), use the **psnp-interval** in router configuration mode.

psnp-interval *milliseconds*

Syntax Description

milliseconds	The interval, specified in milliseconds, at which PSNPs are sent to acknowledge received LSPs. The range is from 1 to 5000 milliseconds. The default value is 50 milliseconds.

Command Default

This command is disabled by default.

Command Modes

Router configuration

Command History

Release	Modification
Release 24.3.1	This command was introduced.

Usage Guidelines

None

Task ID

Task ID	Operation
isis	read, write

Example

This example shows how to configure the PSNP interval to 100 milliseconds for an IS-IS process tagged 1:

```
Router(config)# router isis 1
Router(config-isis)# psnp-interval 100
```

redistribute (IS-IS)

To redistribute routes from one routing protocol into Intermediate System-to-Intermediate System (IS-IS), use the **redistribute** command in address family configuration mode. To remove the **redistribute** command from the configuration file and restore the system to its default condition in which the software does not redistribute routes, use the **no** form of this command.

Border Gateway Protocol (BGP)
redistribute bgp *process-id* [**level-1** | **level-2** | **level-1-2**] [**metric** *metric-value*] [**metric-type** {**internal** | **external**}] [**route-policy** *route-policy-name*]
no redistribute

Connected Routes
redistribute connected [**level-1** | **level-2** | **level-1-2**] [**metric** *metric-value*] [**metric-type** {**internal** | **external**}] [**route-policy** *route-policy-name*]
no redistribute

Intermediate System-to-Intermediate System (IS-IS)
redistribute isis *process-id* [**level-1** | **level-2** | **level-1-2**] [**metric** *metric-value*] [**metric-type** {**internal** | **external**}] [**route-policy** *route-policy-name*] [**down-flag-clear**]
no redistribute

Open Shortest Path First (OSPF)
redistribute ospf *process-id* [**level-1** | **level-2** | **level-1-2**] [**match** {**external** [**1** | **2**] | **internal** | **nssa-external** [**1** | **2**]}] [**metric** *metric-value*] [**metric-type** {**internal** | **external**}] [**route-policy** *route-policy-name*]
no redistribute

Open Shortest Path First Version 3 (OSPFv3)
redistribute ospfv3 *process-id* [**level-1** | **level-2** | **level-1-2**] [**match** {**external** [**1** | **2**] | **internal** | **nssa-external** [**1** | **2**]}] [**metric** *metric-value*] [**metric-type** {**internal** | **external**}] [**route-policy** *route-policy-name*]
no redistribute

Static Routes
redistribute static [**level-1** | **level-2** | **level-1-2**] [**metric** *metric-value*] [**metric-type** {**1** | **2** }] [**route-policy** *route-policy-name*]
no redistribute

Syntax Description

process-id	For the **bgp** keyword, an autonomous system number has the following ranges: • 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. For the **isis** keyword, an IS-IS instance identifier from which routes are to be redistributed. For the **ospf** keyword, an OSPF process name from which routes are to be redistributed. The value takes the form of a string. A decimal number can be entered, but it is stored internally as a string. For the **ospfv3** keyword, an OSPFv3 process name from which routes are to be redistributed. The value takes the form of a string. A decimal number can be entered, but it is stored internally as a string.
level-1	(Optional) Specifies that redistributed routes are advertised in the Level-1 LSP of the router.
level-1-2	(Optional) Specifies that redistributed routes are advertised in the Level-1-2 LSP of the router.
level-2	(Optional) Specifies that redistributed routes are advertised in the Level-2 LSP of the router.
metric *metric-value*	(Optional) Specifies the metric used for the redistributed route. Range is 0 to 16777215. The *metric-value* must be consistent with the IS-IS metric style of the area and topology into which the routes are being redistributed.
metric-type { **internal** | **external** }	(Optional) Specifies the external link type associated with the route advertised into the ISIS routing domain. It can be one of two values: • external • **internal** –Use the **internal** keyword to set IS-IS internal metric-type • **external** –Use the **external** keyword to set IS-IS external metric-type Any route with an internal metric (however large the metric is) is preferred over a route with external metric (however small the metric is).
route-policy *route-policy-name*	(Optional) Specifies the identifier of a configured policy. A policy is used to filter the importation of routes from this source routing protocol to IS-IS.

match { **internal** | **external** [**1** | **2**] | **nsaa-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 interarea 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.
down-flag-clear	(Optional) Specifies that routes redistributed from another IS-IS instance should be advertised with the up/down bit set to zero. This is contrary to the behavior specified by RFC 5305 and RFC 7775 and can lead to route loops.

Command Default

Level 2 is configured if no level is specified.

metric-type: internal

match : If no match keyword is specified, all OSPF routes are redistributed.

Command Modes

Address family configuration

Command History

Release	Modification
Release 3.7.2	This command was introduced.
Release 3.9.0	Asplain format for 4-byte Autonomous system numbers notation was supported.

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.

Note When redistributing routes (into IS-IS) 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.

Use the **redistribute** command to control the redistribution of routes between separate IS-IS instances. To control the propagation of routes between the levels of a single IS-IS instance, use the propagate level, on page 936 command.

Only IPv4 OSPF addresses can be redistributed into IS-IS IPv4 address families.

Task ID

Task ID	Operations
isis	read, write

Examples

In this example, IS-IS instance isp_A readvertises all of the routes of IS-IS instance isp_B in Level 2 LSP. Note that the **level-2** keyword affects which levels instance isp_A advertises the routes in and has no impact on which routes from instance isp_B are advertised. (Any Level 1 routes from IS-IS instance isp_B are included in the redistribution.

```
RP/0/RSP0/CPU0:router(config)# router isis isp_A
RP/0/RSP0/CPU0:router(config-isis)# net 49.1234.2222.2222.2222.00
RP/0/RSP0/CPU0:router(config-isis)# address-family ipv4 unicast
RP/0/RSP0/CPU0:router(config-isis-af)# redistribute isis isp_B level-2
!
RP/0/RSP0/CPU0:router(config)# router isis isp_B
RP/0/RSP0/CPU0:router(config-isis)# is-type level 1
RP/0/RSP0/CPU0:router(config-isis)# net 49.4567.2222.2222.2222.00
RP/0/RSP0/CPU0:router(config-isis)# address-family ipv4 unicast
```

Related Commands

Command	Description
propagate level, on page 936	Propagates routes from one IS-IS level into another level.

remote-psnp-delay

To specify the maximum delay, in milliseconds, that the router should wait for a Partial Sequence Number Protocol (PSNP) acknowledgment from neighbors after sending a Link State Packet (LSP), use the **remote-psnp-delay** command in appropriate mode. To remove the configured maximum delay, use the **no** form of this command.

remote-psnp-delay *milliseconds*

Syntax Description

milliseconds	Maximum delay, in milliseconds, that the router should wait for a PSNP acknowledgment from neighbors after sending a LSP. The range is from 1 through 5000 ms.

Command Default

The command is disabled by default.

Command Modes

IS-IS router lsp-fast-flooding configuration

Interface configuration

Command History

Release	Modification
Release 24.3.1	This command was introduced.

Usage Guidelines

The order of precedence of the configuration using the **remote-psnp-delay** command is as follows:

1. The **remote-psnp-delay** command in interface configuration mode.
2. The value in the PSNP sub-TLV received from the neighbour, if any.
3. The **remote-psnp-delay** command in router configuration mode.

Task ID

Task ID	Operation
isis	read, write

Example

This example shows how to set the PSNP acknowledgment delay to 600 milliseconds on a device.

```
Router# configure terminal
Router(config)# router isis 1
Router(config-isis)# lsp-fast-flooding
Router(config-isis-lsp-fast-flooding)# remote-psnp-delay 600
Router(config-isis-lsp-fast-flooding)# commit
```

This example shows how to set the PSNP acknowledgment delay to 600 milliseconds on an interface.

```
Router# configure terminal
Router(config)# router isis 1
Router(config-isis)# interface GigabitEthernet0/0/0/0
Router(config-isis-if)# remote-psnp-delay 600
```

retransmit-interval (IS-IS)

To configure the amount of time between retransmission of each Intermediate System-to-Intermediate System (IS-IS) link-state packet (LSP) on a point-to-point link, use the **retransmit-interval** command in interface configuration mode. To restore the default value, use the **no** form of this command.

retransmit-interval *seconds* [**level** {**1** | **2**}]
no retransmit-interval [*seconds* [**level** {**1** | **2**}]]

Syntax Description

seconds	Time (in seconds) between consecutive retransmissions of each LSP. It is an integer that should be greater than the expected round-trip delay between any two networking devices on the attached network. Range is 0 to 65535 seconds.
level { **1** | **2** }	(Optional) Specifies routing Level 1 or Level 2 independently.

Command Default

seconds : 5 seconds

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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

The **retransmit-interval** command has no effect on LAN (multipoint) interfaces. On point-to-point links, the value can be increased to enhance network stability.

Because retransmissions occur only when LSPs are dropped, setting this command to a higher value has little effect on reconvergence. The more neighbors networking devices have, and the more paths over which LSPs can be flooded, the higher this value can be made.

Task ID

Task ID	Operations
isis	read, write

Examples

The following example shows how to configure GigabitEthernet interface 0/2/0/1 for retransmission of IS-IS LSPs every 60 seconds for a large serial line:

```
RP/0/RSP0/CPU0:router(config)# router isis isp
RP/0/RSP0/CPU0:router(config-isis)# interface GigabitEthernet 0/2/0/1
RP/0/RSP0/CPU0:router(config-isis-if)# retransmit-interval 60
```

Related Commands

Command	Description
retransmit-throttle-interval, on page 948	Configures the amount of time between retransmissions of any IS-IS LSPs on a point-to-point interface.

retransmit-throttle-interval

To configure minimum interval between retransmissions of different Intermediate System-to-Intermediate System (IS-IS) link-state packets (LSPs) on a point-to-point interface, use the **retransmit-throttle-interval** command in interface configuration mode. To remove the command from the configuration file and restore the system to its default condition, use the **no** form of this command.

retransmit-throttle-interval *milliseconds* [**level** {**1** | **2**}]
no retransmit-throttle-interval [*milliseconds* [**level** {**1** | **2**}]]

Syntax Description

milliseconds	Minimum delay (in milliseconds) between LSP retransmissions on the interface. Range is 0 to 65535.
level { **1** | **2** }	(Optional) Specifies routing Level 1 or Level 2 independently.

Command Default

Default is 0.

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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **retransmit-throttle-interval** command to define the minimum period of time that must elapse between retransmitting any two consecutive LSPs on an interface. The **retransmit-throttle-interval** command may be useful in very large networks with many LSPs and many interfaces as a way of controlling LSP retransmission traffic. This command controls the rate at which LSPs can be re-sent on the interface.

Task ID

Task ID	Operations
isis	read, write

Examples

The following example shows how to configure GigabitEthernet interface 0/2/0/1 to limit the rate of LSP retransmissions to one every 300 milliseconds:

```
RP/0/RSP0/CPU0:router(config)# router isis isp
RP/0/RSP0/CPU0:router(config-isis)# interface GigabitEthernet 0/2/0/1
RP/0/RSP0/CPU0:router(config-isis-if)# retransmit-throttle-interval 300
```

Related Commands

Command	Description
lsp-gen-interval, on page 877	Configures the minimum interval time between regenerating the same LSP.
retransmit-interval (IS-IS), on page 946	Configures the amount of time between retransmission of each IS-IS LSP over a point-to-point link.

router isis

To enable the Intermediate System-to-Intermediate System (IS-IS) routing protocol and to specify an IS-IS instance, use the **router isis** command in global configuration mode. To disable IS-IS routing, use the **no** form of this command.

router isis *instance-id*
no router isis *instance-id*

Syntax Description

instance-id	Name of the routing process. Maximum number of characters is 40.

Command Default

An IS-IS routing protocol is not enabled.

Command Modes

Global configuration

Command History

Release	Modification
Release 3.7.2	This command was introduced.

Usage Guidelines

Use the **router isis** command to create an IS-IS routing process. An appropriate network entity title (NET) must be configured to specify the address of the area (Level 1) and system ID of the router. Routing must be enabled on one or more interfaces before adjacencies may be established and dynamic routing is possible.

Multiple IS-IS processes can be configured. Up to eight processes are configurable. A maximum of five IS-IS instances on a system are supported.

Task ID

Task ID	Operations
isis	read, write

Examples

The following example shows how to configure IS-IS for IP routing:

```
RP/0/RSP0/CPU0:router(config)# router isis isp
RP/0/RSP0/CPU0:router(config-isis)# net 49.0001.0000.0001.00
```

Related Commands

Command	Description
net, on page 922	Configures an IS-IS NET for the routing process.

router isis static neighbor

To enable the Intermediate System-to-Intermediate System (IS-IS) static neighbor and to configure its attributes, use the **router isis static-neighbor** command in Global Configuration mode. To disable IS-IS static neighbor, use the **no** form of this command.

router isis *instance-id* **static-neighbor**

Syntax Description

instance-id	Name of the routing process. Maximum number of characters is 40.

Command Default

None

Command Modes

Global Configuration mode

Command History

Release	Modification
25.2.1	This command was introduced.

Usage Guidelines

You can configure the following attributes by using the **static-neighbor** configuration mode.

system-id	ISIS System-id of the static neighbor.
remote ipv4/ipv6 address	Remote IPv4/IPv6 addresses.
affinity	Affinity link to advertise.
log advertisement changes	Enable logging static neighbor advertisement changes.

Task ID

Task ID	Operations
isis	read, write

Examples

The following example shows how to configure IS-IS static neighbor and its attributes for a router:

```
Router#config
Router(config)#router isis 1
Router(config-isis)#interface GigabitEthernet0/2/0/7
Router(config-isis-if)#static-neighbor
Router(config-isis-static-nbr)#system-id 0001.0002.0007
Router(config-isis-static-nbr)#remote ipv4 address 2.2.2.2
Router(config-isis-static-nbr)#affinity RED
Router(config-isis-static-nbr)#log-advertisement-changes
```

route source first-hop

To replace the originating route with first-hop for multicast traffic, use the **route source first-hop** command in ISIS address-family submode. To remove the first-hop for multicast traffic, use the **no** form of this command.

routesourcefirst-hop

This command has no keywords or arguments.

Command Default

no route source first-hop is enabled.

Command Modes

ISIS address-family submode

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

This command replaces the originating router address with first-hop router address in the RIB table and facilitates computing alternate paths for multicast traffic. This feature is incompatible with other IOS-XR features, such as MPLS-TE inter-area tunnels. You must use the **route source first-hop** command only to support MoFRR with multicast multipath.

Task ID

Task ID	Operations
isis	read, write

Examples

The following example shows how to replace the originating route with first-hop:

```
RP/0/RSP0/CPU0:router(config)# router isis isp
RP/0/RSP0/CPU0:router(config-isis)# address-family ipv4 multicast
RP/0/RSP0/CPU0:router(config-isis-af)# route source first-hop
```

Related Commands

Command	Description
net, on page 922	Configures an IS-IS NET for the routing process.

set default-isis-instance

When multiple IS-IS instances are configured on a router, by default, **show isis** commands display information from all IS-IS instances. To display information from only one specific IS-IS instance, use the **set default-isis-instance** *instance-name* command in the EXEC mode.

When the default IS-IS instance is set, all the show IS-IS commands during the terminal session display only information from that default IS-IS instance. To revert back to the default behaviour of displaying information from all IS-IS instances, use the **set default-isis-instance all** command.

set default-isis-instance { *instance-name* | **all** }

Syntax Description

instance-name	Displays information for the specified *instance-name* when running **show isis** commands.
all	Displays information from all configured IS-IS instances when running **show isis** commands.

Command Default

By default, all configured IS-IS instances are displayed when you run any **show isis** command.

Command Modes

EXEC mode

Command History

Release	Modification
Release 24.2.1	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
isis	read, write

Examples

Example 1: This example shows how to set a specific IS-IS instance as the default instance.

```
RP/0/1/CPU0:ios#set default-isis-instance india
Tue Apr 23 11:47:04.353 CEST
%% Default isis instance is set to  india
```

Example 2: This example uses the **show isis summary** command to display information from the default instance set in Example 1.

```
RP/0/1/CPU0:ios#show isis summary
Tue Apr 23 11:47:10.802 CEST

                                 (L1/L2)     (L1/L2)       (L1/L2)       (IPv4/IPv6)
Tag                    Level     Adj         Routers       LSPs          Routes
---                    -----     ---         -------       ----          ------
```

```
india                  L2      0/1         0/2          0/2          0/0
RP/0/1/CPU0:ios#
```

Example 3: This example shows how to set all configured IS-IS instances as default instances.

```
RP/0/1/CPU0:ios#set default-isis-instance all
Tue Apr 23 11:47:19.308 CEST
%% Default isis instance is set to  all
```

Example 4: This example uses the **show isis summary** command to display information from all the configured instances set in Example 3.

```
RP/0/1/CPU0:ios#show isis summary
Tue Apr 23 11:47:24.697 CEST

                                (L1/L2)     (L1/L2)      (L1/L2)      (IPv4/IPv6)
Tag                    Level    Adj         Routers      LSPs         Routes
---                    -----    ---         -------      ----         ------
eu                     L2       0/2         0/6          0/6          5/5
india                  L2       0/1         0/2          0/2          0/0
RP/0/1/CPU0:ios#
```

set-overload-bit

To configure the router to signal other routers not to use it as an intermediate hop in their shortest path first (SPF) calculations, use the **set-overload-bit** command in router configuration mode. To remove the designation, use the **no** form of this command.

set-overload-bit [**on-startup** {*delay* | **wait-for-bgp**}] [**level** {**1** | **2**}] [**advertise** {**external** | **interlevel**}]
no set-overload-bit [**on-startup** {*delay* | **wait-for-bgp**}] [**level** {**1** | **2**}]

Syntax Description

on-startup	(Optional) Sets the overload bit only temporarily after reboot.
delay	(Optional) Time (in seconds) to advertise when the router is overloaded after reboot. Range is 5 to 86400 seconds (86400 seconds = 1 day).
wait-for-bgp	(Optional) Sets the overload bit on startup until the Border Gateway Protocol (BGP) signals converge or time out.
level { **1** | **2** }	(Optional) Specifies the overload bit for Level 1 or Level 2 independently.

Command Default

The overload bit is not set.

Both Level 1 and Level 2 are configured if no level is specified.

Command Modes

Router 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **set-overload-bit** command to force the router to set the overload bit in its nonpseudonode link-state packets (LSPs). Normally the setting of the overload bit is allowed only when a router experiences problems. For example, when a router is experiencing a memory shortage, the reason might be that the link-state database is not complete, resulting in an incomplete or inaccurate routing table. If the overload bit is set in the LSPs of the unreliable router, other routers can ignore the router in their SPF calculations until it has recovered from its problems. The result is that no paths through the unreliable router are seen by other routers in the Intermediate System-to-Intermediate System (IS-IS) area. However, IP prefixes directly connected to this router are still reachable.

The **set-overload-bit** command can be useful when you want to connect a router to an IS-IS network, but do not want real traffic flowing through it under any circumstances.

Routers with overload bit set are:

- A test router in the lab, connected to a production network.
- A router configured as an LSP flooding server, for example, on a nonbroadcast multiaccess (NBMA) network, in combination with the mesh group feature.

Task ID

Task ID	Operations
isis	read, write

Examples

The following example shows how to configure the overload bit:

```
RP/0/RSP0/CPU0:router(config)# router isis isp
RP/0/RSP0/CPU0:router(config-isis)# set-overload-bit
```

set-attached-bit

To configure an Intermediate System-to-Intermediate System (IS-IS) instance with an attached bit in the Level 1 link-state packet (LSP), use the **set-attached-bit** command in address family configuration mode. To remove the **set-attached-bit** command from the configuration file and restore the system to its default condition, use the **no** form of this command.

set-attached-bit
no set-attached-bit

Command Default

Attached bit is not set in the LSP.

Command Modes

Address family 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **set-attached bit** command to set an IS-IS instance with an attached bit in the Level 1 LSP that allows another IS-IS instance to redistribute Level 2 topology. The attached bit is used when the Level 2 connectivity from another IS-IS instance is advertised by the Level 1 attached bit.

Cisco IOS XR software does not support multiple Level 1 areas in a single IS-IS routing instance. But the equivalent functionality is achieved by redistribution of routes between two IS-IS instances by using the redistribute (IS-IS), on page 940 command.

The attached bit is configured for a specific address family only if the **single-topology** command is not configured.

Note If connectivity for the Level 2 instance is lost, the attached bit in the Level 1 instance LSP continues sending traffic to the Level 2 instance and causes the traffic to be dropped.

Task ID

Task ID	Operations
isis	read, write

Examples

The following example shows how to set the attached bit for a Level 1 instance that allows the Level 2 instance to redistribute routes from the Level 1 instance:

```
RP/0/RSP0/CPU0:router(config)# router isis 1
RP/0/RSP0/CPU0:router(config-isis)# net 49.0001.0001.0001.0001.00
```

```
RP/0/RSP0/CPU0:router(config-isis)# address-family ipv4 unicast
RP/0/RSP0/CPU0:router(config-isis-af)# redistribute isis 2 level 2
!
RP/0/RSP0/CPU0:router(config-isis-af)# interface GigabitEthernet 0/3/0/0
RP/0/RSP0/CPU0:router(config-isis-af-if)# address-family ipv4 unicast
!
!
RP/0/RSP0/CPU0:router(config)# router isis 2
RP/0/RSP0/CPU0:router(config-isis)# is-type level-1
RP/0/RSP0/CPU0:router(config-isis)# net 49.0002.0001.0001.0002.00
RP/0/RSP0/CPU0:router(config-isis)# address-family ipv4 unicast
RP/0/RSP0/CPU0:router(config-isis-af)# set-attached-bit
!
RP/0/RSP0/CPU0:routerfig-isis-af)# interface GigabitEthernet 0/1/0/0
RP/0/RSP0/CPU0:router(config-isis-af-if)# address-family ipv4 unicast
```

Related Commands

Command	Description
redistribute (IS-IS), on page 940	Redistributes routes from one IS-IS instance into another instance.
single-topology, on page 1020	Configures the link topology for IPv4 when IPv6 is configured.

show isis

The **show isis** command displays general information about an IS-IS instance and protocol operation. If the instance ID is not specified, the command shows information about all IS-IS instances.

show isis [**instance** *instance-id*]

Syntax Description

instance *instance-id*	(Optional) Displays the IS-IS adjacencies for the specified IS-IS instance only.

Note
The instance-id argument is the instance identifier (alphanumeric) defined by the **router isis** command.

Command Default

No instance ID specified displays IS-IS adjacencies for all the IS-IS instances.

Command Modes

EXEC

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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

For each instance, the first line of output lists the IS-IS instance ID with the following lines identifying the IS-IS system ID, supported levels (level 1, level 2, or level-1-2), configured area addresses, active area addresses, status (enabled or not) and type (Cisco or IETF) of nonstop forwarding (NSF), and the mode in which the last IS-IS process startup occurred.

Next, the status of each configured address family (or just IPv4 unicast if none are configured) is summarized. For each level (level 1 or level 2), the metric style (narrow or wide) generated and accepted is listed along with the status of incremental shortest path first (iSPF) computation (enabled or not). Then redistributed protocols are listed, followed by the administrative distance applied to the redistributed routes.

Finally, the running state (active, passive, or disabled) and configuration state (active or disabled) of each IS-IS interface is listed.

Task ID

Task ID	Operations
isis	read

Examples

The following is sample output from the **show isis** command:

```
Router# show isis
Wed Aug 20 23:54:55.043 PST DST

IS-IS Router: lab
```

```
System Id: 0000.0000.0002
IS Levels: level-2-only
Manual area address(es):
  49.1122
Routing for area address(es):
  49.1122
Non-stop forwarding: Disabled
Most recent startup mode: Cold Restart
Topologies supported by IS-IS:
  IPv4 Unicast
    Level-2
      Metric style (generate/accept): Narrow/Narrow
      Metric: 10
      ISPF status: Disabled
    No protocols redistributed
    Distance: 115
Interfaces supported by IS-IS:
  Loopback0 is running passively (passive in configuration)
  POS0/1/0/2 is running actively (active in configuration)
  POS0/1/0/3 is running actively (active in configuration
```

This table describes the significant fields shown in the display.

Table 61: show isis Field Descriptions

Field	Description
IS-IS Router	IS-IS instance ID.
System Id	IS-IS system ID.
IS Levels	Supported levels for the instance.
Manual area address(es)	Domain and area.
Routing for area address(es):	Configured area addresses and active area addresses.
Non-stop forwarding	Status (enabled or not) and type (Cisco or IETF) of nonstop forwarding (NSF).
Most recent startup mode	The mode in which the last IS-IS process startup occurred.
Topologies supported by IS-IS	The summary of the status of each configured address family (or just IPv4 unicast if none are configured).
Redistributed protocols	List of redistributed protocols, followed by the administrative distance applied to the redistributed routes.
Metric style (generate/accept)	The status of each configured address family (or just IPv4 unicast if none are configured) is summarized. For each level (level 1 or level 2), the metric style (narrow or wide) generated and accepted is listed along with the status of incremental shortest path first (iSPF) computation (enabled or not).
Interfaces supported by IS-IS	The running state (active, passive, or disabled) and configuration state (active or disabled) of each IS-IS interface.

show isis adjacency

To display Intermediate System-to-Intermediate System (IS-IS) adjacencies, use the **show isis adjacency** command in EXEC mode.

show isis [**instance** *instance-id*] **adjacency** [**level** {**1**|**2**}] [*type interface-path-id*] [**detail**] [**systemid** *system-id*]

Syntax Description

instance *instance-id*	(Optional) Displays the IS-IS adjacencies for the specified IS-IS instance only. • The *instance-id* argument is the instance identifier (alphanumeric) defined by the **router isis** command.
level { **1** | **2** }	(Optional) Displays the IS-IS adjacencies for Level 1 or Level 2 independently.
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.
detail	(Optional) Displays neighbor IP addresses and active topologies.
systemid *system-id*	(Optional) Displays the information for the specified router only.

Command Default

No instance ID specified displays IS-IS adjacencies for all the IS-IS instances.

Both Level 1 and Level 2 are configured if no level is specified.

Command Modes

EXEC

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 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
isis	read

Examples

The following is sample output from the **show isis adjacency** command:

```
RP/0/RSP0/CPU0:router# show isis adjacency

IS-IS p Level-1 adjacencies:
System Id      Interface        SNPA           State Hold     Changed  NSF      BFD
12a4           PO0/1/0/1        *PtoP*         Up    23       00:00:06 Capable  Init
12a4           Gi0/6/0/2        0004.2893.f2f6 Up    56       00:04:01 Capable  Up

Total adjacency count: 2

IS-IS p Level-2 adjacencies:
System Id      Interface        SNPA           State Hold     Changed  NSF      BFD
12a4           PO0/1/0/1        *PtoP*         Up    23       00:00:06 Capable  None
12a4           Gi0/6/0/2        0004.2893.f2f6 Up    26       00:00:13 Capable  Init

Total adjacency count: 2
```

This table describes the significant fields shown in the display.

Table 62: show isis adjacency Field Descriptions

Field	Description
Level-1	Level 1 adjacencies.
Level-2	Level 2 adjacencies.
System ID	Dynamic hostname of the system. The hostname is specified using the **hostname** command. If the dynamic hostname is not known or the **hostname dynamic disable** command has been executed, the 6-octet system ID is used.
Interface	Interface used to reach the neighbor.
SNPA	Data-link address (also known as the Subnetwork Point of Attachment [SNPA]) of the neighbor.
State	Adjacency state of the neighboring interface. Valid states are Down, Init, and Up.
Holdtime	Hold time of the neighbor.
Changed	Time the neighbor has been up (in hours:minutes:seconds).
NSF	Specifies whether the neighbor can adhere to the IETF-NSF restart mechanism.
BFD	Specifies the Bidirectional Forwarding Detection (BFD) status for the interface. Valid status are • None—BFD is not configured. • Init—BFD session is not up. One reason is that other side is not yet enabled. • Up—BFD session has been established. • Down—BFD session holdtime expired.

Related Commands

Command	Description
show isis neighbors, on page 992	Displays information about IS-IS neighbors.

show isis adjacency-log

To display the Intermediate System-to-Intermediate System (IS-IS) adjacency log, use the **show isis adjacency-log** command in EXEC mode.

show isis adjacency-log [**level** {**1** | **2**}] [**last** *number* | **first** *number*]

Syntax Description

level { **1** | **2** }	(Optional) Displays the IS-IS adjacency log for Level 1 or Level 2 independently.
last *number*	(Optional) Specifies that the output is restricted to the last *number* of entries. Range is 1 to 100.
first *number*	(Optional) Specifies that the output is restricted to the first *number* of entries. Range is 1 to 100.

Command Default

No default behavior or values

Command Modes

EXEC

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 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
isis	read

Examples

The following is sample output from the **show isis adjacency-log** command:

```
RP/0/RSP0/CPU0:router# show isis adjacency-log

   IS-IS 10 Level 1 Adjacency log
When           System          Interface          State   Details
4d00h          12a1            PO0/5/0/0          d -> i
4d00h          12a1            PO0/5/0/0          i -> u  New adjacency
                                                          IPv4 Unicast Up
4d00h          12a1            Gi0/6/0/0          d -> u  New adjacency
4d00h          12a1            Gi0/6/0/0          u -> d  Interface state
down
3d17h          12a1            Gi0/6/0/0          d -> u  New adjacency
3d17h          12a1            Gi0/6/0/0          u -> d  Interface state
down
01:44:07       12a1            Gi0/6/0/0          d -> u  New adjacency

   IS-IS 10 Level 2 Adjacency log
When           System          Interface          State   Details
```

```
4d00h          12a1            PO0/5/0/0          d -> i
4d00h          12a1            PO0/5/0/0          i -> u  New adjacency
                                                          IPv4 Unicast Up
4d00h          12a1            Gi0/6/0/0          d -> u  New adjacency
4d00h          12a1            Gi0/6/0/0          u -> d  Interface state
down
3d17h          12a1            Gi0/6/0/0          d -> u  New adjacency
3d17h          12a1            Gi0/6/0/0          u -> d  Interface state
down
01:44:07       12a1            Gi0/6/0/0          d -> u  New adjacency
```

This table describes the significant fields shown in the display.

Table 63: show isis adjacency-log Field Descriptions

Field	Description
When	Elapsed time (in hh:mm:ss) since the event was logged.
System	System ID of the adjacent router.
Interface	Specific interface involved in the adjacency change.
State	State transition for the logged event.
Details	Description of the adjacency change.

show isis checkpoint adjacency

To display the Intermediate System-to-Intermediate System (IS-IS) checkpoint adjacency database, use the **show isis checkpoint adjacency** command in EXEC mode.

show isis [**instance** *instance-id*] **checkpoint adjacency**

Syntax Description

instance *instance-id*	(Optional) Displays the IS-IS checkpoint adjacencies for the specified IS-IS instance only. • The *instance-id* argument is the instance identifier (alphanumeric) defined by the **router isis** command.

Command Default

No instance ID specified displays IS-IS checkpoint adjacencies for all the IS-IS instances.

Command Modes

EXEC

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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **show isis checkpoint adjacency** command to display the checkpointed adjacencies. With this information you can restore the adjacency database during a Cisco proprietary nonstop forwarding (NSF) restart. This command, with the **show isis adjacency** command, can be used to verify the consistency of the two databases.

Task ID

Task ID	Operations
isis	read

Examples

The following is sample output from the **show isis checkpoint adjacency** command:

```
RP/0/RSP0/CPU0:router# show
          isis
          checkpoint
          adjacency

 Interface        Level  System ID        State  Circuit ID          Chkpt ID
 Gi3/0/0/1           1    router-gsr8       Up     0001.0000.0008.04   80011fec
 Gi0/4/0/1           1    router-gsr9       Up     0001.0000.0006.01   80011fd8
 Gi/0/0/1           2    router-gsr8      Up     0001.0000.0008.04   80011fc4
```

This table describes the significant fields shown in the display.

Table 64: show isis checkpoint adjacency Field Descriptions

Field	Description
Interface	Interface used to reach the neighbor.
Level	Lists either routers with Level 1 or Level 2 adjacency configured.
System ID	Dynamic hostname of the system. The hostname is specified using the **hostname** command. If the dynamic hostname is not known or **hostname dynamic disable** command has been executed, the 6-octet system ID is used.
State	State of the neighboring interface.
Circuit ID	Unique ID issued to a circuit at its creation.
Chkpt ID	Unique ID issued to the checkpoint at its creation.

Related Commands

Command	Description
show isis adjacency, on page 961	Displays IS-IS adjacencies.
show isis checkpoint lsp, on page 969	Displays the IS-IS checkpoint LSP database.

show isis checkpoint interface

To display the Intermediate System-to-Intermediate System (IS-IS) checkpoint interfaces, use the **show isis checkpoint interface** command in EXEC mode.

show isis checkpoint interface

This command has no keywords or arguments.

Command Default

No default behavior or values

Command Modes

EXEXEC EC

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 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
isis	read

Examples

The following is sample output from the **show isis checkpoint interface** command:

```
RP/0/RSP0/CPU0:router# show isis checkpoint interface

IS-IS 10 checkpoint interface
Interface         Index  CircNum  DIS Areas    Chkpt ID
PO0/5/0/0         0      0        NONE         80002fe8
Gi0/6/0/0         1      3        L1L2         80002fd0
```

This table describes the significant fields shown in the display.

Table 65: show isis checkpoint interface Field Descriptions

Field	Description
Interface	Interface used to reach the neighbor.
Index	Interface index assigned to an interface upon its creation.
CircNum	Unique ID issued to a circuit internally.
DIS Areas	Designated Intermediate System area.

Field	Description
Chkpt ID	Unique ID issued to the checkpoint at its creation.

show isis checkpoint lsp

To display the Intermediate System-to-Intermediate System (IS-IS) checkpoint link-state packet (LSP) protocol data unit (PDU) identifier database, use the **show isis checkpoint lsp** command in EXEC mode.

show isis [**instance** *instance-id*] **checkpoint lsp**

Syntax Description

instance *instance-id*	(Optional) Displays the IS-IS checkpoint LSPs for the specified instance only. • The *instance-id* argument is the instance identifier (alphanumeric) defined by the **router isis** command.

Command Default

No instance ID specified displays IS-IS checkpoint LSPs for all the IS-IS instances.

Command Modes

EXEC

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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

The checkpointed LSPs displayed by this command are used to restore the LSP database during a Cisco-proprietary nonstop forwarding (NSF) restart. The **show isis checkpoint lsp** command, with the **show isis database** command, may be used to verify the consistency of the two databases.

Task ID

Task ID	Operations
isis	read

Examples

The following is sample output from the **show isis checkpoint lsp** command:

```
RP/0/RSP0/CPU0:router#
            show isis checkpoint lsp


Level  LSPID              Chkpt ID
1      router-gsr6.00-00  80011f9c
1      router-gsr6.01-00  80011f88
1      router-gsr8.00-00  80011f74
1      router-gsr9.00-00  80011f60
2      router-gsr6.00-00  80011f4c
2      router-gsr6.01-00  80011f38
2      router-gsr8.00-00  80011f24
2      router-gsr9.00-00  80011f10
Total LSP count: 8 (L1: 4, L2 4, local L1: 2, local L2 2)
```

This table describes the significant fields shown in the display.

Table 66: show isis checkpoint lsp Field Descriptions

Field	Description
Level	Routers with Level 1 or Level 2 adjacency configured.
LSPID	LSP identifier. The first six octets form the system ID of the router that originated the LSP. The next octet is the pseudonode ID. When this byte is 0 , the LSP describes links from the system. When it is nonzero, the LSP is a so-called nonpseudonode LSP. This is similar to a router link-state advertisement (LSA) in the Open Shortest Path First (OSPF) protocol. The LSP describes the state of the originating router. For each LAN, the designated router for that LAN creates and floods a pseudonode LSP, describing all systems attached to that LAN. The last octet is the LSP number. If there is more data than can fit in a single LSP, the LSP is divided into multiple LSP fragments. Each fragment has a different LSP number. An asterisk (*) indicates that the LSP was originated by the system on which this command is issued.
Chkpt ID	Unique ID issued to the checkpoint at its creation.

Related Commands

Command	Description
show isis checkpoint adjacency, on page 965	Displays the IS-IS checkpoint adjacency database.
show isis database, on page 971	Displays the IS-IS link-state database.

show isis database

To display the Intermediate System-to-Intermediate System (IS-IS) link-state packet (LSP) database, use the **show isis database** command in EXEC mode.

show isis [**instance** *instance-id*] **database** [**level** {**1** | **2**}] [**update**] [**summary**] [**detail**] [**verbose**] [***lsp-id***]

Syntax Description

instance *instance-id*	(Optional) Displays the IS-IS LSP database for the specified instance only. • The *instance-id* argument is the instance identifier (alphanumeric) defined by the **router isis** command.
level { **1** | **2** }	(Optional) Displays the IS-IS LSP database for Level 1 or Level 2 independently.
update	(Optional) Displays contents of LSP database managed by update thread.
summary	(Optional) Displays the LSP ID number, sequence number, checksum, hold time, and bit information.
detail	(Optional) Displays the contents of each LSP.
verbose	(Optional) Displays the contents of each LSP.
* | *lsp-id*	(Optional) LSP protocol data units (PDUs) identifier. Displays the contents of a single LSP by its ID number or may contain an * as a wildcard character.

Command Default

No instance ID specified displays the IS-IS LSP database for all the IS-IS instances.

Both Level 1 and Level 2 is configured if no level is specified.

Command Modes

EXEC

Command History

Release	Modification
Release 3.7.2	This command was introduced.

Usage Guidelines

Each of the options for the **show isis database** command can be entered in an arbitrary string within the same command entry. For example, the following are both valid command specifications and provide the same output: **show isis database detail level 2** and **show isis database level 2 detail** .

The **summary** keyword used with this command allows you to filter through a large IS-IS database and quickly identify problematic areas.

Task ID

Task ID	Operations
isis	read

show isis database-log

To display the entries in the Intermediate System-to-Intermediate System (IS-IS) database log, use the **show isis database-log** command in EXEC mode.

show isis database-log [**level** {**1** | **2**}] [**last** *number* | **first** *number*]

Syntax Description

level { **1** | **2** }	(Optional) Displays the database log for Level 1 or Level 2 independently.
last *number*	(Optional) Specifies that the output be restricted to the last *number* of entries. Range is 1 to 1000.
first *number*	(Optional) Specifies that the output be restricted to the first *number* of entries. Range is 1 to 1000.

Command Default

Both Level 1 and Level 2 are configured if no level is specified.

Command Modes

EXEC

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 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
isis	read

Examples

The following is sample output from the **show isis database-log** command:

```
RP/0/RSP0/CPU0:router# show isis database-log

   IS-IS 10 Level 1 Link State Database Log
                                 New LSP                 Old LSP
WHEN     LSPID                Op  Seq Num    Holdtime OL  Seq Num    Holdtime OL
01:17:19 12b1.03-00           REP 0x00000003 1200      0  0x00000002 340       0
001:06:20 12b1.00-00          REP 0x000001d8 1200      0  0x000001d7 375       0
01:06:00 12b1.03-00           REP 0x00000004 1200      0  0x00000003 520       0
01:05:46 12a1.00-00           REP 0x000001fc 1200      0  0x000001fb 425       0
00:55:01 12b1.00-00           REP 0x000001d9 1200      0  0x000001d8 520       0
00:53:39 12b1.03-00           REP 0x00000005 1200      0  0x00000004 459       0
00:53:19 12a1.00-00           REP 0x000001fd 1200      0  0x000001fc 453       0
00:42:12 12b1.00-00           REP 0x000001da 1200      0  0x000001d9 431       0
00:39:56 12b1.03-00           REP 0x00000006 1200      0  0x00000005 376       0
00:38:54 12a1.00-00           REP 0x000001fe 1200      0  0x000001fd 334       0
00:29:10 12b1.00-00           REP 0x000001db 1200      0  0x000001da 418       0
00:27:22 12b1.03-00           REP 0x00000007 1200      0  0x00000006 446       0
```

```
00:25:10 12a1.00-00           REP 0x000001ff 1200      0  0x000001fe 375       0
00:17:04 12b1.00-00           REP 0x000001dc 1200      0  0x000001db 473
```

This table describes the significant fields shown in the display.

Table 67: show isis database-log Field Descriptions

Field	Description
WHEN	Elapsed time (in hh:mm:ss) since the event was logged.
LSPID	LSP identifier. The first six octets form the system ID of the router that originated the LSP. The next octet is the pseudonode ID. When this byte is 0 , the LSP describes links from the system. When it is nonzero, the LSP is a so-called nonpseudonode LSP. This is similar to a router link-state advertisement (LSA) in the Open Shortest Path First (OSPF) protocol. The LSP describes the state of the originating router. For each LAN, the designated router for that LAN creates and floods a pseudonode LSP, describing all systems attached to that LAN. The last octet is the LSP number. If there is more data than can fit in a single LSP, the LSP is divided into multiple LSP fragments. Each fragment has a different LSP number. An asterisk (*) indicates that the LSP was originated by the system on which this command is issued.
New LSP	New router or pseudonode appearing in the topology.
Old LSP	Old router or pseudonode leaving the topology.
Op	Operation on the database: inserted (INS) or replaced (REP).
Seq Num	Sequence number for the LSP that allows other systems to determine if they have received the latest information from the source.
Holdtime	Time the LSP remains valid (in seconds). An LSP hold time of 0 indicates that this LSP was purged and is being removed from the link-state database (LSDB) of all routers. The value indicates how long the purged LSP stays in the LSDB before being completely removed.
OL	Overload bit. Determines if the IS is congested. If the Overload bit is set, other routers do not use this system as a transit router when calculating routers. Only packets for destinations directly connected to the overloaded router are sent to this router.

Related Commands

Command	Description
show isis database, on page 971	Displays the IS-IS link-state packet (LSP) database.

show isis fast-reroute

To display per-prefix LFA information, use the **show isis fast-reroute** command in EXEC mode.

show isis fast-reroute
A.B.C.D/length | **detail** | **summary** | **sr-only**

Syntax Description

A.B.C.D/length	Network to show per-prefix LFA information.
detail	Use to display tiebreaker information about the backup.
summary	Use to display the number of prefixes having protection per priority.
sr-only	Use to display SR-labeled prefixes only.

Command Default

None

Command History

Release	Modification
Release 4.0.1	This command was introduced.
Release 6.3.2	The **sr-only** keyword was added.

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
isis	read

The following is sample output from **show isis fast-reroute** command that displays per-prefix LFA information:

```
RP/0/RSP0/CPU0:router# show isis fast-reroute 10.1.6.0/24

L1 10.1.6.0/24 [20/115]
     via 10.3.7.47, POS0/3/0/1, router2
       FRR backup via 10.1.7.145, GigabitEthernet0/1/0/3, router3
```

The following is sample output from **show isis fast-reroute detail** command that displays tie-breaker information about the backup:

```
RP/0/RSP0/CPU0:router# show isis fast-reroute 10.1.6.0/24 detail

L1 10.1.6.0/24 [20/115] low priority
     via 10.3.7.47, POS0/3/0/1, router2
       FRR backup via 10.1.7.145, GigabitEthernet0/1/0/3, router3
       P: No, TM: 30, LC: Yes, NP: No, D: No
     src router2.00-00, 192.168.0.47
```

```
   L2 adv [20] native, propagated
```

The following is sample output from **show isis fast-reroute summary** command that displays the number of prefixes having protection per priority:

```
RP/0/RSP0/CPU0:router#show isis fast-reroute summary
IS-IS frr IPv4 Unicast FRR summary

                                    Critical   High       Medium     Low        Total

                                    Priority   Priority   Priority      Priority

Prefixes reachable in L1
  All paths protected     0              0             2                8               10

  Some paths protected 0               0             1                3               4

  Unprotected                0              0             1                3
 4
  Protection coverage    0.00%     0.00%     75.00%     78.57%     77.78%
Prefixes reachable in L2
  All paths protected     0              0             0                0               0

  Some paths protected 0               0             1                0               1

  Unprotected                0              0             0                0
 0
  Protection coverage    0.00%     0.00%     100.00%   0.00%      100.00%
```

The following is sample output from **show isis fast-reroute sr-only** command that displays fast-reroute repair paths for prefixes associated with a segment routing prefix SID:

```
RP/0/RSP0/CPU0:router#show isis fast-reroute sr-only
IS-IS 1 IPv4 Unicast FRR backups

Codes: L1 - level 1, L2 - level 2, ia - interarea (leaked into level 1)
       df - level 1 default (closest attached router), su - summary null
       C - connected, S - static, R - RIP, B - BGP, O - OSPF
       E - EIGRP, A - access/subscriber, M - mobile, a - application
       i - IS-IS (redistributed from another instance)
       D - Downstream, LC - Line card disjoint, NP - Node protecting
       P - Primary path, SRLG - SRLG disjoint, TM - Total metric via backup

Maximum parallel path count: 8

L2 20.1.0.101/32 [10/115]
     via 10.1.1.101, GigabitEthernet0/0/0/2, r101, SRGB Base: 16000, Weight: 0
        Backup path: TI-LFA (link), via 10.4.1.103, GigabitEthernet0/0/0/1 r103, SRGB Base:
 16000, Weight: 0
          P node: r103.00 [20.1.0.103], Label: ImpNull
          Q node: r102.00 [20.1.0.102], Label: 24001
          Prefix label: 16101
          Backup-src: r101.00
L2 20.1.0.102/32 [30/115]
     via 10.1.1.101, GigabitEthernet0/0/0/2, r101, SRGB Base: 16000, Weight: 0
        Backup path: TI-LFA (link), via 10.4.1.103, GigabitEthernet0/0/0/1 r103, SRGB Base:
 16000, Weight: 0
          P node: r103.00 [20.1.0.103], Label: ImpNull
          Q node: r102.00 [20.1.0.102], Label: 24001
          Prefix label: ImpNull
          Backup-src: r102.00
L2 20.1.0.103/32 [20/115]
```

```
    via 10.4.1.103, GigabitEthernet0/0/0/1, r103, SRGB Base: 16000, Weight: 0
       Backup path: TI-LFA (link), via 10.1.1.101, GigabitEthernet0/0/0/2 r101, SRGB Base:
16000, Weight: 0
          P node: r102.00 [20.1.0.102], Label: 16102
          Q node: r103.00 [20.1.0.103], Label: 24001
          Prefix label: ImpNull
          Backup-src: r103.00
```

show isis hostname

To display the entries in the Intermediate System-to-Intermediate System (IS-IS) router name-to-system ID mapping table, use the **show isis hostname** command in EXEC mode.

show isis [**instance** *instance-id*] **hostname**

Syntax Description

instance *instance-id*	(Optional) Displays the IS-IS router name-to-system ID mapping table for the specified IS-IS instance only. The *instance-id* argument is the instance identifier (alphanumeric) defined by the **router isis** command.

Command Default

No instance ID specified displays the IS-IS router name-to-system ID mapping table for all the IS-IS instances.

Command Modes

EXEC

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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

The **show isis hostname** command does not display entries if the dynamic hostnames are disabled.

Task ID

Task ID	Operations
isis	read

Examples

The following is sample output from the **show isis hostname** command with the **instance** and *instance-id* values specified:

```
RP/0/RSP0/CPU0:router# show isis instance isp hostname

ISIS isp hostnames
    Level  System ID      Dynamic Hostname
     1     0001.0000.0005 router
     2   * 0001.0000.0011 router-11
```

This table describes the significant fields shown in the display.

Table 68: show isis instance isp hostname Field Descriptions

Field	Description
Level	IS-IS level of the router.
System ID	Dynamic hostname of the system. The hostname is specified using the **hostname** command. If the dynamic hostname is not known or **hostname dynamic disable** command has been executed, the 6-octet system ID is used.
Dynamic Hostname	Hostname of the router.
*	Local router.

Related Commands

Command	Description
hostname	Specifies the name of the local router.
hostname dynamic disable, on page 858	Enables the IS-IS routing protocol to dynamically update the mapping of router names to system IDs.

show isis interface

To display information about the Intermediate System-to-Intermediate System (IS-IS) interfaces, use the **show isis interface** command in EXEC mode.

show isis interface [*type interface-path-id* | **level** {**1** | **2**}] [**brief**]

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.
level { **1** | **2** }	(Optional) Displays IS-IS interface information for Level 1 or Level 2 independently.
brief	(Optional) Displays brief interface output.

Command Default

Displays all IS-IS interfaces.

Command Modes

EXEC

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 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
isis	read

Examples

The following is sample output from the **show isis interface** command:

```
RP/0/RSP0/CPU0:router#show isis interface
          GigabitEthernet 0/3/0/2
 Gi /3/0/2                  Enabled
  Adjacency Formation:      Enabled
  Prefix Advertisement:     Enabled
  BFD:                      Disabled
  BFD Min Interval:         150
  BFD Multiplier:           3
```

```
Circuit Type:             level-2-only
Media Type:               P2P
Circuit Number:           0
Extended Circuit Number:  67111168
Next P2P IIH in:          4 s
LSP Rexmit Queue Size:    0

Level-2
  Adjacency Count:        1
  LSP Pacing Interval:    33 ms
  PSNP Entry Queue Size:  0

CLNS I/O
  Protocol State:         Up
  MTU:                    4469

IPv4 Unicast Topology:    Enabled
  Adjacency Formation:    Running
  Prefix Advertisement:   Running
  Metric (L1/L2):         10/100
  MPLS LDP Sync (L1/L2):  Disabled/Disabled
IPv6 Unicast Topology:    Disabled (Not cfg on the intf)

IPv4 Address Family:      Enabled
  Protocol State:         Up
  Forwarding Address(es): 10.3.10.143
  Global Prefix(es):      10.3.10.0/24
IPv6 Address Family:      Disabled (No topology enabled which uses IPv6)

LSP transmit timer expires in 0 ms
LSP transmission is idle
Can send up to 9 back-to-back LSPs in the next 0 ms
```

This table describes the significant fields shown in the display.

Table 69: show isis interface Field Descriptions

Field	Description
GigabitEthernet0/6/0/0	Status of the interface, either enabled or disabled.
Adjacency formation	Status of adjacency formation, either enabled or disabled.
Prefix Advertisement	Status of advertising connected prefixes, either enabled or disabled.
BFD	Status of Bidirectional Forwarding Detection (BFD), either enabled or disabled.
BFD Min Interval	BFD minimum interval.
BFD Multiplier	BFD multiplier.
Circuit Type	Levels the interface is running on (circuit-type configuration) which may be a subset of levels on the router.
Media Type	Media type on which IS-IS is running.
Circuit Number	Unique ID assigned to a circuit internally (8-bit integer).
Extended Circuit Number	Valid only for point-to-point interfaces (32-bit integer).

Field	Description
LSP Rexmit Queue Size	Number of LSPs pending retransmission on the interface.
Adjacency Count	Number of adjacencies formed with a neighboring router that supports the same set of protocols.
PSNP Entry Queue Size	Number of SNP entries pending inclusion in the next PSNP.
LAN ID	ID of the LAN.
Priority (Local/DIS)	Priority of this interface or priority of the Designated Intermediate System.
Next LAN IIH in	Time (in seconds) in which the next LAN hello message is sent.
LSP Pacing Interval	Interval at which the link-state packet (LSP) transmission rate (and by implication the reception rate of other systems) is to be reduced.
Protocol State	Running state of the protocol (up or down).
MTU	Link maximum transmission unit (MTU).
SNPA	Data-link address (also known as the Subnetwork Point of Attachment [SNPA]) of the neighbor.
All Level-n ISs	Status of interface membership in Layer 2 multicast group. The status options are Yes or reason for not being a member of the multicast group.
IPv4 Unicast Topology	Status of the topology, either enabled or disabled.
Adjacency Formation	Status of adjacency formation. The status options are Running or a reason for not being ready to form adjacencies.
Prefix Advertisement	Status of advertising prefixes, either enabled or disabled.
Metric (L1/L2)	IS-IS metric for the cost of the adjacency between the originating router and the advertised neighbor, or the metric of the cost to get from the advertising router to the advertised destination (which can be an IP address, an end system (ES), or a connectionless network service (CLNS) prefix).
MPLS LDP Sync (L1/L2)	Status of LDP IS-IS synchronization, either enabled or disabled. When enabled, the state of synchronization (Sync Status) is additionally displayed as either achieved or not achieved.
IPv4 Address Family	Status of the address family, either enabled or disabled.
Protocol State	State of the protocol.
Forwarding Address(es)	Addresses on this interface used by the neighbor for next-hop forwarding.
Global Prefix(es)	Prefixes for this interface included in the LSP.
LSP transmit timer expires in	LSP transmission expiration timer interval (in milliseconds).

Field	Description
LSP transmission is	State of LSP transmission. Valid states are • idle • in progress • requested • requested and in progress

The following is sample output from the **show isis interface** command with the **brief** keyword:

```
RP/0/0/CPU0:router# show isis interface brief

    Interface      All     Adjs    Adj Topos  Adv Topos  CLNS   MTU    Prio
                   OK    L1   L2    Run/Cfg    Run/Cfg                L1   L2
-----------------  ---  ---------  ---------  ---------  ----  ----  --------
PO0/5/0/0          Yes    1    1      1/1        1/1      Up    4469    -    -
Gi0/6/0/0          Yes    1*   1*     1/1        1/1      Up    1497   64   64
```

This table describes the significant fields shown in the display.

Table 70: show isis interface brief Field Descriptions

Field	Description
Interface	Name of the interface.
All OK	Everything is working as expected for this interface.
Adjs L1 L2	Number of L1 and L2 adjacencies over this interface.
Adj Topos Run/Cfg	Number of topologies that participate in forming adjacencies. Number of topologies that were configured to participate in forming adjacencies.
Adv Topos Run/Cfg	Number of topologies that participate in advertising prefixes. Number of topologies that were configured to participate in advertising prefixes.
CLNS	Status of the Connectionless Network Service. Status options are Up or Down.
MTU	Maximum transfer unit size for the interface.
Prio L1 L2	Interface L1 priority. Interface L2 priority.

show isis lsp-log

To display link-state packet (LSP) log information, use the **show isis lsp-log** command in EXEC mode.

show isis [**instance** *instance-id*] **lsp-log** [**level** {**1** | **2**}] [**last** *number* | **first** *number*]

Syntax Description

instance *instance-id*	(Optional) Displays the LSP log information for the specified IS-IS instance only. • The *instance-id* argument is the instance identifier (alphanumeric) defined by the **router isis** command.
level { **1** | **2** }	(Optional) Displays the Intermediate System-to-Intermediate System (IS-IS) link-state database for Level 1 or Level 2 independently.
last *number*	(Optional) Specifies that the output be restricted to the last *number* of entries. Range is 1 to 20.
first *number*	(Optional) Specifies that the output be restricted to the first *number* of entries. Range is 1 to 20.

Command Default

No instance ID specified displays the LSP log information for all the IS-IS instances.

Both Level 1 and Level 2 are configured if no level is specified.

Command Modes

EXEC

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 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
isis	read

Examples

The following is sample output from the **show isis lsp-log** command with the **instance** and *instance-id* values specified:

```
RP/0/RSP0/CPU0:router# show isis instance isp lsp-log

ISIS isp Level 1 LSP log
  When       Count      Interface       Triggers
00:02:36         1
00:02:31         1                      LSPREGEN
00:02:26         1      PO4/1           DELADJ
00:02:24         1      PO4/1           NEWADJ
```

```
00:02:23         1      Gi5/0           DIS
00:01:27         1      Lo0             IPDOWN
00:01:12         1      Lo0             IPUP


ISIS isp Level 2 LSP log
  When       Count      Interface       Triggers
00:02:36         1
00:02:30         1                      LSPREGEN
00:02:26         1      PO4/1           DELADJ
00:02:24         1      PO4/1           NEWADJ
00:02:23         1      Gi5/0           DIS
00:02:21         1                      AREASET
00:01:27         1      Lo0             IPDOWN
00:01:12         1      Lo0             IPUP
```

This table describes the significant fields shown in the display.

Table 71: show isis instance isp lsp-log Field Descriptions

Field	Description
Level	IS-IS level of the router.
When	How long ago (in hh:mm:ss) an LSP rebuild occurred. The last 20 occurrences are logged.
Count	Number of events that triggered this LSP run. When there is a topology change, often multiple LSPs are received in a short period. A router waits 5 seconds before running a full LSP, so it can include all new information. This count denotes the number of events (such as receiving new LSPs) that occurred while the router was waiting its 5 seconds before running full LSP.
Interface	Interface that corresponds to the triggered reasons for the LSP rebuild.
Triggers	A list of all reasons that triggered an LSP rebuild. The triggers are • AREASET—area set changed • ATTACHFLAG—bit attached • CLEAR— **clear** command • CONFIG—configuration change • DELADJ—adjacency deleted • DIS—DIS changed • IFDOWN—interface down • IPADDRCHG—IP address change • IPDEFORIG—IP def-orig • IPDOWN—connected IP down • IFDOWN—interface down • IPEXT—external IP • IPIA—nterarea IP • IPUP—connected IP up • LSPDBOL—LSPDBOL bit • LSPREGEN—LSP regeneration • NEWADJ— new adjacency

show isis mesh-group

To display Intermediate System-to-Intermediate System (IS-IS) mesh group information, use the **show isis mesh-group** command in EXEC mode.

show isis [**instance** *instance-id*] **mesh-group**

Syntax Description

instance *instance-id*	(Optional) Displays the mesh group information for the specified IS-IS instance only. • The *instance-id* argument is the instance identifier (alphanumeric) defined by the **router isis** command.

Command Default

No instance ID specified displays the IS-IS mesh group information for all the IS-IS instances.

Command Modes

EXEC

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 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
isis	read

Examples

The following is sample output from the **show isis mesh-group** command with the **instance** and *instance-id* values specified:

```
RP/0/RSP0/CPU0:router# show isis instance isp mesh-group

 ISIS isp Mesh Groups

 Mesh group 6:
 GigabitEthernet 0/4/0/1
```

This table describes the significant fields shown in the display.

Table 72: show isis instance isp mesh-group Field Descriptions

Field	Description
Mesh group	Mesh group number to which this interface is a member. A mesh group optimizes link-state packet (LSP) flooding in nonbroadcast multiaccess (NBMA) networks with highly meshed, point-to-point topologies. LSPs that are first received on interfaces that are part of a mesh group are flooded to all interfaces except those in the same mesh group.
GigabitEthernet0/4/0/1	Interface belonging to mesh group 6.

show isis mpls traffic-eng adjacency-log

To display a log of Multiprotocol Label Switching traffic engineering (MPLS TE) adjacency changes for an Intermediate System-to-Intermediate System (IS-IS) instance, use the **show isis mpls traffic-eng adjacency-log** command in EXEC mode.

show isis [**instance** *instance-id*] **mpls traffic-eng adjacency-log** [**last** *number* | **first** *number*]

Syntax Description

instance *instance-id*	(Optional) Displays the MPLS TE adjacency changes for the specified IS-IS instance only. • The *instance-id* argument is the instance identifier (alphanumeric) defined by the **router isis** command.
last *number*	(Optional) Specifies that the output is restricted to last *number* of entries. Range is 1 to 20.
first *number*	(Optional) Specifies that the output is restricted to first *number* of entries. Range is 1 to 20.

Command Default

No instance ID specified displays MPLS TE adjacency changes for all the IS-IS instances.

Command Modes

EXEC

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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **show isis mpls traffic-eng adjacency-log** command to display the status of MPLS TE adjacencies.

Task ID

Task ID	Operations
isis	read

Examples

The following is sample output from the **show isis mpls traffic-eng adjacency-log** command with the **instance** and *instance-id* values specified:

```
RP/0/RSP0/CPU0:router# show isis instance isp mpls traffic-eng adjacency-log

IS-IS isp Level-2 MPLS Traffic Engineering adjacency log
When      Neighbor ID        IP Address       Interface Status
00:03:36  router-6           172.17.1.6         PO0/3/0/1 Up
00:03:36  router-6           172.17.1.6         PO0/3/0/1 Down
00:02:38  router-6           172.17.1.6         PO0/3/0/1 Up
```

This table describes the significant fields shown in the display.

Table 73: show isis instance isp mpls traffic-eng adjacency-log Field Descriptions

Field	Description
When	Time (in hh:mm:ss) since the entry was recorded in the log.
Neighbor ID	Identification value of the neighbor.
IP Address	Neighbor IP Version 4 (IPv4) address.
Interface	Interface from which a neighbor is learned.
Status	Up (active) or Down (disconnected).

Related Commands

Command	Description
show isis mpls traffic-eng advertisements, on page 989	Displays the last flooded record from MPLS traffic engineering.

show isis mpls traffic-eng advertisements

To display the latest flooded record from Multiprotocol Label Switching traffic engineering (MPLS TE) for an Intermediate System-to-Intermediate System (IS-IS) instance, use the **show isis mpls traffic-eng advertisements** command in EXEC mode.

show isis [**instance** *instance-id*] **mpls traffic-eng advertisements**

Syntax Description

instance	*instance-id*	(Optional) Displays the latest flooded record from MPLS TE for the specified IS-IS instance only. • The *instance-id* argument is the instance identifier (alphanumeric) defined by the **router isis** command.

Command Default

No instance ID specified displays the latest flooded record from MPLS TE for all the IS-IS instances.

Command Modes

EXEC

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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **show isis mpls traffic-eng advertisements** command to verify that MPLS TE is flooding its record and that the bandwidths are correct.

Task ID

Task ID	Operations
isis	read

Examples

The following is sample output from the **show isis mpls traffic-eng advertisements** command with the **instance** and *instance-id* values specified:

```
RP/0/RSP0/CPU0:router# show isis instance isp mpls traffic-eng advertisements

 ISIS isp Level-2 MPLS Traffic Engineering advertisements
       System ID: router-9
       Router ID: 172.18.0.9
       Link Count: 1
         Link[0]
           Neighbor System ID: router-gsr6 (P2P link)
           Interface IP address: 172.18.0.9
           Neighbor IP Address: 172.18.0.6
           Admin. Weight: 0
           Physical BW: 155520000 bits/sec
```

```
Reservable BW global: 10000000 bits/sec
Reservable BW sub: 0 bits/sec
Global pool BW unreserved:
  [0]: 10000000 bits/sec, [1]: 10000000 bits/sec
  [2]: 10000000 bits/sec, [3]: 10000000 bits/sec
  [4]: 10000000 bits/sec, [5]: 10000000 bits/sec
  [6]: 10000000 bits/sec, [7]: 10000000 bits/sec
Sub pool BW unreserved:
  [0]: 0 bits/sec, [1]: 0 bits/sec
  [2]: 0 bits/sec, [3]: 0 bits/sec
  [4]: 0 bits/sec, [5]: 0 bits/sec
  [6]: 0 bits/sec, [7]: 0 bits/sec
Affinity Bits: 0x00000000
```

This table describes the significant fields shown in the display.

Table 74: show isis instance isp mpls traffic-eng advertisements Field Descriptions

Field	Description
System ID	Dynamic hostname of the system. The hostname is specified using the **hostname** command. If the dynamic hostname is not known or if the **hostname dynamic disable** command has been executed, the 6-octet system ID is used.
Router ID	MPLS TE router ID.
Link Count	Number of links that MPLS TE advertised.
Neighbor System ID	System ID of a neighbor number in an area. The six bytes directly preceding the n-selector are the system ID. The system ID length is a fixed size and cannot be changed. The system ID must be unique throughout each area (Level 1) and throughout the backbone (Level 2). In an IS-IS routing domain, each router is represented by a 6-byte hexadecimal system ID. When network administrators maintain and troubleshoot networking devices, they must know the router name and corresponding system ID.
Interface IP address	IP address of the interface.
Neighbor IP Address	IP address of the neighbor.
Admin. Weight	Administrative weight associated with this link.
Physical BW	Link bandwidth capacity (in bits per second).
Reservable BW	Reservable bandwidth on this link.
Global pool BW unreserved	Unreserved bandwidth that is available in the global pool.
Sub pool BW unreserved	Amount of unreserved bandwidth that is available in the subpool.
Affinity Bits	Link attribute flags being flooded. Bits are MPLS-TE specific.

Related Commands

Command	Description
show isis mpls traffic-eng adjacency-log, on page 987	Displays a log of MPLS TE adjacency changes for IS-IS.

show isis neighbors

To display information about Intermediate System-to-Intermediate System (IS-IS) neighbors, use the **show isis neighbors** command in EXEC mode.

show isis [**instance** *instance-id*] **neighbors** [*type interface-path-id* | **summary**] [**detail**] [**systemid** *system-id*]

Syntax Description

instance *instance-id*	(Optional) Displays the IS-IS neighbor information for the specified IS-IS instance only. • The *instance-id* argument is the instance identifier (alphanumeric) defined by the **router isis** command.
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.
summary	(Optional) Displays neighbor status count for each level.
detail	(Optional) Displays additional details.
systemid *system-id*	(Optional) Displays the information for the specified neighbor only.

Command Default

No instance ID specified displays neighbor information for all the IS-IS instances.

Both Level 1 and Level 2 are configured if no level is specified.

Command Modes

EXEC

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 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
isis	read

Examples

The following is sample output from the **show isis neighbors** command with the **instance** and *instance-id* values specified:

```
Total neighbor count: 3
RP/0/RSP0/CPU0:router# show isis instance isp neighbors detail

IS-IS isp neighbors:
System Id      Interface        SNPA           State Holdtime Type IETF-NSF
e222e          Gi0/1/0/0        *PtoP*         Up    23       L1   Capable
  Area Address(es): 00
  IPv4 Address(es): 10.1.0.45*
  IPv6 Address(es): fe80::212:daff:fe6b:68a8*
    Topologies: 'IPv4 Unicast'
  Uptime: 01:09:44
  IPFRR: LFA Neighbor: elise
         LFA IPv4 address: 10.100.1.2
         LFA Router address: 192.168.0.45
e333e Gi0/1/0/0.1      0012.da6b.68a8 Up    8        L1   Capable
e333e          Gi0/1/0/0.1      0012.da6b.68a8 Up    8        L1   Capable
  Area Address(es): 00
  IPv4 Address(es): 10.100.1.2*
  Topologies: 'IPv4 Unicast'
  Uptime: 01:09:46
  IPFRR: LFA Neighbor: elise
         LFA IPv4 address: 10.1.0.45
         LFA Router address: 192.168.0.45
         LFA Interface: Gi0/1/0/0
m44i           Gi0/1/0/1        0012.da62.e0a8 Up    7        L1   Capable
  Area Address(es): 00 11
  IPv4 Address(es): 10.1.2.47*
  IPv6 Address(es): fe80::212:daff:fe62:e0a8*
    Topologies: 'IPv4 Unicast'
  Uptime: 01:09:33

Total neighbor count: 3
```

This table describes the significant fields shown in the display.

Table 75: show isis instance isp neighbors Field Descriptions

Field	Description
System ID	Dynamic hostname of the system. The hostname is specified using the **hostname** command. If the dynamic hostname is not known or **hostname dynamic disable** command has been executed, the 6-octet system ID is used.
Interface	Interface through which the neighbor is reachable.
SNPA	Data-link address (also known as the Subnetwork Point of Attachment [SNPA]) of the neighbor.
State	Adjacency state of the neighboring interface. Valid states are: Down, Init, and Up.
Holdtime	Hold time of the neighbor.
Type	Type of adjacency.

Field	Description
IETF-NSF	Specifies whether the neighbor can adhere to the IETF-NSF restart mechanism. Valid states are Capable and Unable.
Area Address(es)	Number of area addresses on this router.
IPv4 Address(es)	IPv4 addresses configured on this router.
Topologies	Address and subaddress families for which IS-IS is configured.
Uptime	Time (in hh:mm:ss) that the neighbor has been up.
IPFRR: LFA Neighbor	IP fast reroute (IPFRR) loop-free alternate (LFA) neighbor.
LFA IPv4 address:	Address of the LFA.
LFA Interface:	LFA interface.

The following is sample output from the **show isis neighbors** command with the **summary** keyword specified:

```
RP/0/RSP0/CPU0:router# show isis instance isp neighbors summary

ISIS isp neighbor summary:
    State          L1       L2     L1L2
    Up              0        0        2
    Init            0        0        0
    Failed          0        0        0
```

This table describes the significant fields shown in the display.

Table 76: show isis neighbors summary Field Descriptions

Field	Description
State	State of the neighbor is up, initialized, or failed.
L1	Number of Level 1 neighbors.
L2	Number of Level 2 neighbors.
L1L2	Number of Level 1 and 2 neighbors.

Related Commands

Command	Description
show isis adjacency, on page 961	Displays IS-IS adjacencies.

show isis protocol

To display summary information about an Intermediate System-to-Intermediate System (IS-IS) instance, use the **show isis protocol** command in EXEC mode.

show isis [**instance** *instance-id*] **protocol**

Syntax Description

instance *instance-id*	(Optional) Displays the IS-IS adjacencies for the specified IS-IS instance only. • The *instance-id* argument is the instance identifier (alphanumeric) defined by the **router isis** command.

Command Default

No instance ID specified displays IS-IS adjacencies for all the IS-IS instances.

Command Modes

EXEC

Command History

Release	Modification
Release 3.7.2	This command was introduced.
Release 24.2.1	This command was was modified. The fields to indicate the memory state of the IS-IS protocol were included.

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
isis	read

Examples

The following is sample output from the **show isis protocol** command:

```
Router# show isis protocol

IS-IS Router: isp
      System Id: 0001.0000.0011
      IS Levels: level-1-2
      Manual area address(es):
        49

      Routing for area address(es):
        49
      Non-stop forwarding: Cisco Proprietary NSF Restart enabled
      Process startup mode: Cold Restart
      Topologies supported by IS-IS:
        IPv4 Unicast
          Level-1 iSPF status: Dormant (awaiting initial convergence)
              Level-2 iSPF status: Dormant (awaiting initial convergence)
          No protocols redistributed
```

```
            Distance: 115
          Interfaces supported by IS-IS:
            Loopback0 is running passively (passive in configuration)
            GigabitEthernet 0/4/0/1 is running actively (active in configuration)
            GigabitEthernet 0/5/0/1 is running actively (active in configuration)
Memory state: Normal
     Last change: 2024 Jan 27 23:05:48.952
     Last memory usage/limit: 608032/1610612400 0%
   Memory state monitoring interval: 60 s
   Memory usage/limit: 19983152/1610612400 1%
   Memory usage specifics (count/bytes):
     Nodes: 3/1104     /* Nodes are using 1,104 bytes across 3 nodes.*/
     LSPs: 3/1248      /* LSPs are using 1,248 bytes across 3 LSPs.*/
     Links: 12/14592   /* Links are using 14,592 bytes across 12 links.*/
     Prefixes: 12/3264 /* Prefixes are using 3,264 bytes across 12 prefixes. */
     Nets: 8/4640      /* Nets are using 4,640 bytes across 8 nets.*/
    Paths: 0/0        /* Paths are currently not using any memory (0 bytes across 0 paths).
 */
```

This table describes the significant fields shown in the display.

Table 77: show isis protocol Field Descriptions

Field	Description
System ID:	Dynamic hostname of the system. The hostname is specified using the **hostname** command. If the dynamic hostname is not known or **hostname dynamic disable** command has been executed, the 6-octet system ID is used.
IS Levels:	IS-IS level of the router.
Manual area address(es)	Area addresses that are manually configured.
Routing for areaaddress(es)	Area addresses for which this router provides the routing.
Non-stop forwarding:	Status and name of nonstop forwarding (NSF).
Process startup mode:	Mode in which the last process startup occurred. Valid modes are: • Cisco Proprietary NSF Restart • IETF NSF Restart • Cold Restart

Field	Description
iSPF status:	State of incremental shortest path first (iSPF) configuration for this IS-IS instance. Four states exist: Disabled if iSPF has not been configured but is awaiting a full SPF to compile the topology for use by the iSPF algorithm. Dormant if iSPF has been configured but is awaiting initial convergence before initializing. Awake if iSPF has been configured but is awaiting a full SPF to compile the topology for use by the iSPF algorithm. Active if IS-IS is ready to consider using the iSPF algorithm whenever a new route calculation needs to be run.
No protocols redistributed:	No redistributed protocol information exists to be displayed.
Distance:	Administrative distance for this protocol.
Memory state	Current status of memory usage for IS-IS process.
Last change	Timestamp of the last change in memory state.
Last memory usage/limit	The previous recorded memory usage in bytes and the maximum memory limit. The percentage shows the proportion of the limit used.
Memory state monitoring interval	The frequency at which the memory state is checked for the IS-IS process.
Memory usage/limit	The current memory usage against the maximum memory limit for the IS-IS process. The percentage shows the current proportion of the limit used.
Memory usage specifics (count/bytes)	A detailed breakdown of memory usage by various components or functions within the IS-IS process, indicating the count and memory consumption.
Nodes	Number of nodes and memory in bytes used by the nodes.
LSPs	Number of nodes and memory in bytes used by the LSPs.
Links	Number of nodes and memory in bytes used by the links.
Prefixes	Number of nodes and memory in bytes used by the prefixes.
Nets	Number of nodes and memory in bytes used by the network entities.
Paths	Number of nodes and memory in bytes used by the paths.

show isis route

To display IP reachability information for an Intermediate System-to-Intermediate System (IS-IS) instance, use the **show isis route** command in EXEC mode.

show isis [**instance** *instance-id*] [**ipv4** | **ipv6** | **afi-all**] [**unicast** | **multicast** [**topology** {**all***topo-name*}] | **safi-all**] **route** [*ip-address mask* | *ip-address*/*length* [**longer-prefixes**]] [**summary**] **[multicast-intact]** [**backup**] [**detail**] [**sr-only**]

Syntax Description

instance *instance-id*	(Optional) Displays the IP reachability information for the specified IS-IS instance only. • The *instance-id* argument is the instance identifier (alphanumeric) defined by the **router isis** command.
ipv4	(Optional) Specifies IP Version 4 address prefixes.
ipv6	(Optional) Specifies IP Version 6 address prefixes.
afi-all	(Optional) Specifies all address prefixes.
unicast	(Optional) Specifies unicast address prefixes.
multicast	(Optional) Specifies multicast address prefixes.
topology	(Optional) Specifies IS-IS paths to intermediate systems.
all	(Optional) Specifies all topologies.
topology *topo-name*	(Optional) Specifies topology table information and name of the topology table.
safi-all	(Optional) Specifies all secondary address prefixes.
ip-address	(Optional) Network IP address about which routing information should be displayed.
mask	(Optional) Network mask specified in either of two ways: • Network mask can be a four-part, dotted decimal address. For example, 255.0.0.0 indicates that each bit equal to 1 means the corresponding address bit is a network address. • Network mask can be indicated as a slash (/) and number. For example, /8 indicates that the first 8 bits of the mask are ones, and the corresponding bits of the address are the network address.
/ length	(Optional) Length of the IP 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. Range is 0 to 32.
longer-prefixes	(Optional) Displays route and more-specific routes.
summary	(Optional) Displays topology summary information.

multicast-intact	(Optional) Displays multicast intact information for this entry.
systemid	(Optional) Displays multicast information by system ID.
backup	(Optional) Displays information about ISIS routes not selected for installation in the RIB.
detail	(Optional) Displays link-state packet (LSP) details.
sr-only	(Optional) Displays SR-labeled prefixes only.

Command Default

No instance ID specified displays the IP reachability information for all the IS-IS instances.

Command Modes

EXEC

Command History

Release	Modification
Release 3.7.2	This command was introduced.
Release 3.9.0	Support for IPv6 was added.
Release 6.3.2	The **sr-only** keyword was added.

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
isis	read

Examples

The following is sample output from the **show isis route** command:

```
RP/0/RSP0/CPU0:router# show isis route

IS-IS isp IPv4 Unicast routes
Codes: L1 - level 1, L2 - level 2, ia - interarea (leaked into level 1)
df - level 1 default (closest attached router), su - summary null
C - connected, S - static, R - RIP, B - BGP, O - OSPF
i - IS-IS (redistributed from another instance)

Maximum parallel path count: 8

L2 10.76.240.6/32 [4/115]
via 10.76.245.252, SRP0/1/0/2, isp2
via 10.76.246.252, SRP0/1/0/0, isp2
C 10.76.240.7/32
is directly connected, Loopback0
L2 10.76.240.9/32 [256/115]
via 10.76.249.2, GigabitEthernet 0/3/0/0, isp3
L2 10.76.240.10/32 [296/115]
via 10.76.249.2, GigabitEthernet 0/3/0/0, isp3
C 10.76.245.0/24
```

```
  is directly connected, SRP0/1/0/2
  C 10.76.246.0/24
  is directly connected, SRP0/1/0/0
  C 10.76.249.0/26
  is directly connected, GigabitEthernet 0/3/0/0
  L2 10.101.10.0/24 [296/115]
  via 10.76.249.2, GigabitEthernet 0/3/0/0, isp3
```

This table describes the significant fields shown in the display.

Table 78: show isis route ipv4 unicast Field Descriptions

Field	Description
C172.18.0.0/24	Connected route for GigabitEthernet interface 0/5/0/0.
C 172.19.1.0/24	Connected route for GigabitEthernet interface 0/4/0/1.
L1 172.35.0.0/24 [10]	Level 1 route to network 172.35.0.0/24.
C 172.18.0/24	Connected route for loopback interface 0.

This is sample output from the **show isis route** command with **sr-only** keyword that shows only routes associated with a segment routing prefix SID:

```
RP/0/RSP0/CPU0:router# show isis route sr-only

IS-IS 1 IPv4 Unicast routes

Codes: L1 - level 1, L2 - level 2, ia - interarea (leaked into level 1)
       df - level 1 default (closest attached router), su - summary null
       C - connected, S - static, R - RIP, B - BGP, O - OSPF
       E - EIGRP, A - access/subscriber, M - mobile, a - application
       i - IS-IS (redistributed from another instance)

Maximum parallel path count: 8

C  20.1.0.100/32
     is directly connected, Loopback0
L2 20.1.0.101/32 [10/115]
     via 10.1.1.101, GigabitEthernet0/0/0/2, r101, SRGB Base: 16000, Weight: 0
L2 20.1.0.102/32 [30/115]
     via 10.1.1.101, GigabitEthernet0/0/0/2, r101, SRGB Base: 16000, Weight: 0
L2 20.1.0.103/32 [20/115]
     via 10.4.1.103, GigabitEthernet0/0/0/1, r103, SRGB Base: 16000, Weight: 0
```

show isis spf-log

To display how often and why the router has run a full shortest path first (SPF) calculation, use the **show isis spf-log** command in EXEC mode.

show isis [**instance** *instance-id*] [[**ipv4** | **ipv6** | **afi-all**] [**unicast** | **multicast** [**topology** { **all** *topo-name* }] | **safi-all**]] **spf-log** [**level** { **1** | **2** }] [**ispf** | **fspf** | **prc** | **nhc**] [**detail** | **verbose** | **plfrr** | **ppfrr**] [**last** *number* | **first** *number*]

Syntax Description

instance *instance-id*	(Optional) Displays the IS-IS SPF log for the specified IS-IS instance only.
ipv4	(Optional) Specifies IP Version 4 address prefixes.
ipv6	(Optional) Specifies IP Version 6 address prefixes.
afi-all	(Optional) Specifies all address prefixes.
unicast	(Optional) Specifies unicast address prefixes.
multicast	(Optional) Specifies multicast address prefixes.
topology all | *topo-name*	(Optional) Specifies topology table information for all topologies or for the specified topology table (*top-name*).
safi-all	(Optional) Specifies all secondary address prefixes.
level { **1** | **2** }	(Optional) Displays the IS-IS SPF log for Level 1 or Level 2 independently.
ispf	(Optional) Specifies incremental SPF entries only.
fspf	(Optional) Specifies full SPF entries only.
prc	(Optional) Specifies partial route calculations only.
nhc	(Optional) Specifies next-hop route calculations only.
detail	(Optional) Specifies detailed output. Includes a breakdown of the time taken to perform the calculation and changes resulting from the calculation.
verbose	(Optional) Specifies verbose output.
last *number*	(Optional) Specifies that the output is restricted to the last *number* of entries. Range is 1 to 210.
first *number*	(Optional) Specifies that the output is restricted to the first *number* of entries. Range is 1 to 210.

Command Default

No instance ID specified displays IS-IS adjacencies for all the IS-IS instances.

Both Level 1 and Level 2 are configured if no level is specified.

Displays all types of route calculation (not just fspf, ispf, and prc).

Command Modes

EXEC

Command History

Release	Modification
Release 3.7.2	This command was introduced.
Release 3.9.0	Support for IPv6 was added.

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
isis	read

Examples

The following is sample output from the **show isis spf-log** command:

```
Router# show isis spf-log

      IS-IS 1 Level 1 IPv4 Unicast Route Calculation Log
                  Time  Total Trig
 Timestamp   Type (ms)  Nodes Count First Trigger LSP Triggers
 ----------- ---- ----  ----- ----- ----- ------- --- --------
 --- Thurs Aug 19 2004 ---
 12:00:50.787  FSPF  1    1     3   ensoft-grs7.00-00 LSPHEADER TLVCODE
 12:00:52.846  FSPF  1    1     1   ensoft-grs7.00-00 LSPHEADER
 12:00:56.049  FSPF  1    1     1   ensoft-grs7.00-00 TLVCODE
 12:01:02.620  FSPF  1    1     2   ensoft-grs7.00-00 NEWADJ LINKTLV

      IS-IS 1 Level 1 IPv4 Unicast Route Calculation Log
                  Time  Total Trig
 Timestamp   Type (ms)  Nodes Count First Trigger LSP Triggers
 ----------- ---- ----  ----- ----- ----- ------- --- --------
 --- Mon Aug 19 2004 ---
 12:00:50.790  FSPF  0    1     4   ensoft-grs7.00-00 LSPHEADER TLVCODE
 12:00:54.043  FSPF  1    1     2   ensoft-grs7.00-00 NEWADJ LSPHEADER
 12:00:55.922  FSPF  1    2     1   ensoft-grs7.00-00 NEWLSPO
 12:00:56.724  FSPF  1    13    1   ensoft-grs7.00-00 NEWLSPO
```

This table describes the significant fields shown in the display.

Table 79: show isis spf-log ipv4 unicast Field Descriptions

Field	Description
Level	IS-IS level of the router.
Timestamp	Time when the SPF calculation started.
Duration	Number of milliseconds taken to complete this SPF run. Elapsed time is wall clock time, not CPU time.

Field	Description
Nodes	Number of routers and pseudonodes (LANs) that make up the topology calculated in this SPF run.
Trig Count	Number of events that triggered this SPF run. When there is a topology change, often multiple link-state packets (LSPs) are received in a short time. Depending on the configuration of the **spf-interval** command, a router may wait for a fixed period of time before running a router calculation. This count denotes the number of triggering events that occurred while the router was waiting to run the calculation. For a full description of the triggering events, see *List of Triggers*.
First Trigger LSP	LSP ID stored by the router whenever a full SPF calculation is triggered by the arrival of a new LSP. The LSP ID can suggest the source of routing instability in an area. If multiple LSPs are causing an SPF run, only the LSP ID of the first received LSP is remembered.
Triggers	List of all reasons that triggered a full SPF calculation. For a list of possible triggers, see *List of Triggers*.

This table lists triggers of a full SPF calculation.

Table 80: List of Triggers

Trigger	Description
PERIODIC	Runs a full SPF calculation very 15 minutes.
NEWLEVEL	Configured new level (using is-type) on this router.
RTCLEARED	Cleared IS-IS topology on the router.
MAXPATHCHANGE	Changed IP maximum parallel path.
NEWMETRIC	Changed link metric.
ATTACHFLAG	Changed Level 2 Attach bit.
ADMINDIST	Configured another administrative distance for the IS-IS instance on this router.
NEWADJ	Created a new adjacency to another router.
DELADJ	Deleted adjacency.
BACKUP	Installed backup route.
SEEDISPF	Seed incremental SPF.
NEXTHOP	Changed IP next-hop address.
NEWLSP0	New LSP 0 appeared in the topology.
LSPEXPIRED	Some LSP in the link-state database (LSDB) has expired.
LSPHEADER	Changed important LSP header fields.

Trigger	Description
TLVCODE	Type, length, and value (TLV) objects code mismatch, indicating that different TLV objects are included in the newest version of an LSP.
LINKTV	Changed Link TLV content.
PREFIXTLV	Changed Prefix TLV content.
AREAADDRTLV	Changed Area address TLV content.
IP ADDRTLV	Changed IP address TLV content.
TUNNEL	Changed RRR tunnel.

The following is sample output from the **show isis spf-log** command with the **first** keyword specified:

```
Router# show isis spf-log first 2

    IISIS isp Level 1 IPv4 Unicast Route Calculation Log
                 Time  Total Trig
Timestamp   Type (ms)  Nodes Count First Trigger LSP   Triggers
  Mon Aug 16 2004
19:25:35.140  FSPF  1    1     1             12a5.00-00 NEWLSP0
19:25:35.646  FSPF  1    1     1                         NEWADJ


   IISIS isp Level 2  IPv4 Unicast Route Calculation Log
                 Time  Total Trig
Timestamp   Type (ms)  Nodes Count First Trigger LSP   Triggers
  Mon Aug 16 2004
19:25:35.139  FSPF  1    1     1             12a5.00-00 NEWLSP0
19:25:35.347  FSPF  1    1     2             12a5.00-00 NEWSADJ TLVCODE
```

This table describes the significant fields shown in the display.

Table 81: show isis spf-log first Field Descriptions

Field	Description
Level	IS-IS level of the router.
Timestamp	Time at which the SPF calculation started.
Type	Type of route calculation. The possible types are incremental SPF (iSPF),full SPF (FSPF), or partial route calculation (PRC).
Time (ms)	Number of milliseconds taken to complete this SPF run. Elapsed time is wall clock time, not CPU time.
Nodes	Number of routers and pseudonodes (LANs) that make up the topology calculated in this SPF run.

Field	Description
Trig Count	Number of events that triggered this SPF run. When there is a topology change, often multiple link-state packets (LSPs) are received in a short time. Depending on the configuration of the **spf-interval** command, a router may wait for a fixed period of time before running a router calculation. This count denotes the number of triggering events that occurred while the router was waiting to run the calculation. For a full description of the triggering events, see *List of Triggers*.
First Trigger LSP	LSP ID stored by the router whenever a full SPF calculation is triggered by the arrival of a new LSP. The LSP ID can suggest the source of routing instability in an area. If multiple LSPs are causing an SPF run, only the LSP ID of the first received LSP is remembered.
Triggers	List of all reasons that triggered a full SPF calculation. For a list of possible triggers, see *List of Triggers*.

The following is sample output from the **show isis spf-log** command with the **detail** keyword specified:

```
Router# show isis spf-log detail

    IISIS isp Level 1 IPv4 Unicast Route Calculation Log
                Time  Total Trig
Timestamp   Type (ms)  Nodes Count First Trigger LSP   Triggers
  Mon Aug 16 2004
19:25:35.140  FSPF  1    1     1             12a5.00-00 NEWLSP0
  Delay:              51ms (since first trigger)
  SPT Calculation
    CPU Time:         0ms
    Real Time:        0ms
  Prefix Updates
    CPU Time:         1ms
    Real Time:        1ms
  New LSP Arrivals:    0
  Next Wait Interval: 200ms
                            Results
                       Reach Unreach Total
   Nodes:                  1       0     1
   Prefixes (Items)
     Critical Priority:    0       0     0
     High Priority:        0       0     0
     Medium Priority       0       0     0
     Low Priority          0       0     0

     All Priorities        0       0     0
   Prefixes (Routes)
     Critical Priority:    0       -     0
     High Priority:        0       -     0
     Medium Priority       0        -    0
     Low Priority:         0        -    0

     All Priorities        0        -    0
```

This table describes the significant fields shown in the display.

Table 82: show isis spf-log detail Field Descriptions

Field	Description
Level	IS-IS level of the router.
Timestamp	Time at which the SPF calculation started.
Type	Type of route calculation. The possible types are incremental ispf, full SPF (FSPF), or partial route calculation (PRC).
Time (ms)	Number of milliseconds taken to complete this SPF run. Elapsed time is wall clock time, not CPU time.
Nodes	Number of routers and pseudonodes (LANs) that make up the topology calculated in this SPF run.
Trig Count	Number of events that triggered this SPF run. When there is a topology change, often multiple link-state packets (LSPs) are received in a short time. Depending on the configuration of the **spf-interval** command, a router may wait for a fixed period of time before running a router calculation. This count denotes the number of triggering events that occurred while the router was waiting to run the calculation. For a full description of the triggering events, see *List of Triggers*.
First Trigger LSP	LSP ID stored by the router whenever a full SPF calculation is triggered by the arrival of a new LSP. The LSP ID can suggest the source of routing instability in an area. If multiple LSPs are causing an SPF run, only the LSP ID of the first received LSP is remembered.
Triggers	List of all reasons that triggered a full SPF calculation. For a list of possible triggers, see *List of Triggers*.
Delay	Two different delays exist: **1.** The delay between the time when the route calculation was first triggered and the time when it was run. **2.** The delay between the end of the last route calculation and the start of this one. This is used to verify that the SPF-interval timers are working correctly, and is only reported for calculations after the first delay.
CPU Time	Two different CPU times exist: **1.** CPU time (in milliseconds) taken to calculate the shortest path tree (SPT). **2.** CPU time (in milliseconds) taken to perform the prefix updates.
Real Time	Two different real times exist: **1.** Real time (in milliseconds) taken to calculate the shortest path tree (SPT). **2.** Real time (in milliseconds) taken to perform the prefix updates.
New LSP Arrivals	Number of LSP arrivals since the start of this route calculation.

Field	Description
Next Wait Interval	Enforced delay until the next route calculation can be run, based on the **spf-interval** command configuration.
Reach	Number of reachable nodes or prefixes.
Unreach	Number of unreachable nodes or prefixes.
Total	Total number of nodes or prefixes at various priorities.

Related Commands

Command	Description
spf-interval, on page 1022	Sets IS-IS throttling of shortest path first (SPF) calculations.

show isis statistics

To display Intermediate System-to-Intermediate System (IS-IS) traffic counters, use the **show isis statistics** command in EXEC mode.

show isis [**instance** *instance-id*] **statistics** [*type interface-path-id*]

Syntax Description

instance *instance-id*	(Optional) Displays the IS-IS traffic statistics for the specified IS-IS instance only. • The *instance-id* argument is the instance identifier (alphanumeric) defined by the **router isis** command.
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

No instance ID specified displays IS-IS traffic statistics for all the IS-IS instances.

IS-IS traffic statistics are displayed for all interfaces.

Command Modes

EXEC

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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

The **show isis statistics** command displays IS-IS traffic counters for the specified interface or all traffic counters if no interface is specified.

Task ID

Task ID	Operations
isis	read

Examples

The following is sample output from the **show isis statistics** command that shows all traffic counters:

```
Router#show isis statistics
IS-IS isp statistics:
  Fast PSNP cache (hits/tries): 164115/301454
```

```
Fast CSNP cache (hits/tries): 41828/43302
Fast CSNP cache updates: 2750
LSP checksum errors received: 0
LSP Dropped: 1441
SNP Dropped: 1958
UPD Max Queue size: 2431
Average transmit times and rate:
  Hello:          0 s,     987947 ns,          4/s
  CSNP:           0 s,    1452987 ns,          0/s
  PSNP:           0 s,    1331690 ns,          0/s
  LSP:            0 s,    1530018 ns,          1/s
Average process times and rate:
  Hello:          0 s,     874584 ns,         41/s
  CSNP:           0 s,     917925 ns,         29/s
  PSNP:           0 s,    1405458 ns,          0/s
  LSP:            0 s,    4352850 ns,          0/s
Level-1:
  LSPs sourced (new/refresh): 3376/2754
  IPv4 Unicast
    SPF calculations           : 520
    ISPF calculations          : 0
    Next-hop Calculations      : 0
    Partial Route Calculations : 0
Level-2:
  LSPs sourced (new/refresh): 4255/3332
  IPv4 Unicast
    SPF calculations           : 432
    ISPF calculations          : 0
    Next-hop Calculations      : 8
    Partial Route Calculations : 0
Interface GigabitEthernet0/1/0/1.1:
Level-1 Hellos (sent/rcvd): 22398/25633
Level-1 DR Elections      : 66
Level-1 LSPs (sent/rcvd)  : 246/7077
Level-1 CSNPs (sent/rcvd) : 0/33269
Level-1 PSNPs (sent/rcvd) : 22/0
Level-1 LSP Flooding Duplicates  : 25129
Level-2 Hellos (sent/rcvd): 22393/67043
Level-2 DR Elections      : 55
Level-2 LSPs (sent/rcvd)  : 265/437
Level-2 CSNPs (sent/rcvd) : 0/86750
Level-2 PSNPs (sent/rcvd) : 0/0
Level-2 LSP Flooding Duplicates  : 78690
```

This table describes the significant fields shown in the display.

Table 83: show isis statistics Field Descriptions

Field	Description
Fast PSNP cache (hits/tries)	Number of successful lookups (hits) along with the number of lookup attempts (tries). To save time or processing power when receiving multiple copies of the same LSP, IS-IS attempts to look up incoming LSPs to see if they have been received recently.
Fast CSNP cache (hits/tries)	Number of successful lookups (hits) along with the number of lookup attempts (tries). To reduce CSNP construction time, IS-IS maintains a cache of CSNPs and attempts to look up CSNP in this cache before transmission on the interface.

Field	Description
Fast CSNP cache updates	Number of times the CSNP cache has been updated since the last clearing of statistics. The cache is updated on LSP addition or removal from the database.
LSP checksum errors received	Number of internal checksum errors received in LSPs.
IIH (LSP/SNP) dropped	Number of hello, LSP, and SNP messages dropped.
IIH (UPD) Max Queue size	Maximum number of queued packets.
Average transmit times and rate	Average time taken to transmit the pdu type across all interfaces and the corresponding rate at which the pdu type is being transmitted.
Average process times and rate	Average time taken to process an incoming pdu type across all interfaces and the corresponding rate at which the pdu type is being received.
LSPs sourced (new/refresh)	Number of LSPs this IS-IS instance has created or refreshed. To find more details on these LSPs, use the **show isis lsp-log** command.
SPF calculations	Number of shortest path first (SPF) calculations. SPF calculations are performed only when the topology changes. They are not performed when external routes change. The interval at which SPF calculations are performed is configured using the **spf-interval** command.
iSPF calculations	Number of incremental shortest path first (iSPF) calculations. iSPF calculations are performed only when ISPF has been configured in the isis address family configuration submode.
Partial Route Calculations	Number of partial route calculations (PRCs). PRCs are processor intensive. Therefore, it may be useful to limit their number, especially how often a PRC is done, especially on slower networking devices. Increasing the PRC interval reduces the processor load on the router, but might slow the rate of convergence. The interval at which PRC calculations are performed is configured using the **spf-interval** command.
Level-(1/2) (LSPs/CSNPs/PSNPs/Hellos) (sent/rcvd)	Number of LSPs, Complete Sequence Number Packets (CSNPs), Partial Sequence Number Packets (PSNPs), and hello packets sent or received on this interface.
PTP Hellos (sent/rcvd)	Point-to-point (PTP) hellos sent and received.
LSP Retransmissions	Total number of retransmissions on each IS-IS LSP on a point-to-point interface. The LSP retransmission interval can be configured using the **retransmit-throttle-interval** command.
Level-(1.2) DRElections	Total number of Designated Intermediate System elections that have taken place. These counts are maintained on an individual level basis.
LSP Flooding Duplicates	Number of duplicate LSPs filtered from flooding to the neighbor. In case of parallel interfaces to the same neighbor, IS-IS optimizes the flooding by avoiding sending the same LSP copy on other interfaces.

Related Commands

Command	Description
show isis spf-log, on page 1001	Displays how often and why the router has run a full SPF calculation.
spf-interval, on page 1022	Sets IS-IS throttling of shortest path first (SPF) calculations.

show isis topology

To display a list of connected Intermediate System-to-Intermediate System (IS-IS) routers in all areas, use the **show isis topology** command in EXEC mode.

show isis [**instance instance-id**] [[**ipv4** | **ipv6** | **afi-all**] [**unicast** | **multicast** [**topology** {**all** | **topo-name**}] | **safi-all**]] | **summary** | **level** {**1** | **2**} **[multicast-intact]** [**systemid system-id**] [**detail**]

Syntax Description

instance *instance-id*	(Optional) Displays the IS-IS topology for the specified IS-IS instance only. • The *instance-id* argument is the instance identifier (alphanumeric) defined by the **router isis** command.
ipv4	(Optional) Specifies IP Version 4 address prefixes.
ipv6	(Optional) Specifies IP Version 6 address prefixes.
afi-all	(Optional) Specifies all address prefixes.
unicast	(Optional) Specifies unicast address prefixes.
multicast	(Optional) Specifies multicast address prefixes.
topology *topo-name*	(Optional) Specifies topology table information and name of the topology table.
safi-all	(Optional) Specifies all secondary address prefixes.
summary	(Optional) Displays a brief list of the IS-IS topology.
level { **1** | **2** }	(Optional) Displays the IS-IS link-state topology for Level 1 or Level 2 independently.
multicast-intact	(Optional) Displays multicast intact information on the IS-IS topology.
systemid *system-id*	(Optional) Displays the information for the specified router only.
detail	(Optional) Displays detailed information on the IS-IS topology.

Command Default

No instance ID specified displays a list of connected routers in all areas for all the IS-IS instances.

Both Level 1 and Level 2 is configured if no level is specified.

Command Modes

EXEC

Command History

Release	Modification
Release 3.7.2	This command was introduced.
Release 3.9.0	Support for IPv6 was added.

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.

Use the **show isis topology** command to verify the presence and connectivity among all routers in all areas.

Use the **show isis topology flex-algo dataplane** command to display information on Flexible Algorithms.

Task ID

Task ID	Operations
isis	read

Examples

The following is sample output from the **show isis topology** command:

```
RP/0/RSP0/CPU0:router# show isis topology

IS-IS isp paths to (Level-1) routers
    System Id       Metric  Next-hop Interface       SNPA
    ensoft-5        10      ensoft-5        PO0/4/0/1       *PtoP*
    ensoft-5        10      ensoft-5        Gi0/5/0/0       0003.6cff.0680
    ensoft-11       --

    IS-IS isp paths to (Level-2) routers
    System Id       Metric  Next-hop Interface       SNPA
    ensoft-5        10      ensoft-5        PO0/4/0/1       *PtoP*
    ensoft-5        10      ensoft-5        Gi0/5/0/0       0003.6cff.0680
    ensoft-11       --
```

This table describes the significant fields shown in the display.

Table 84: show isis topology ipv4 unicast Field Descriptions

Field	Description
System ID	Dynamic hostname of the system. The hostname is specified using the **hostname** command. If the dynamic hostname is not known or **hostname dynamic disable** command has been executed, the 6-octet system ID is used.
Metric	Metric assigned to the link and used to calculate the cost from each router using the links in the network to other destinations. Range is 1 to 16777214. Default is 1 to 63 for narrow metric and 1 to 16777214 for wide metric. 0 is set internally if no metric has been specified by the user.
Next-hop	Address of the next-hop.
Interface	Interface used to reach the neighbor.
SNPA	Data-link address (also known as the Subnetwork Point of Attachment [SNPA]) of the neighbor.

The following is sample output from the **show isis topology** command with the **summary** keyword specified:

```
RP/0/RSP0/CPU0:router# show isis topology summary
```

```
IS-IS 10 IS Topology Summary IPv4 Unicast
                        L1                          L2
               Reach  UnReach  Total      Reach  UnReach  Total
              ------- ------- -------    ------- ------- -------
Router nodes:        1       1       2          1       1       2
Pseudo nodes:        0       0       0          0       0       0

 Total nodes:        1       1       2          1       1       2
```

This table describes the significant fields shown in the display.

Table 85: show isis topology summary Field Descriptions

Field	Description
L1/L2	IS-IS level of the router.
Reach	Number of router nodes or pseudonodes that are reachable.
UnReach	Number of router nodes or pseudonodes that are unreachable.
Total	Total number of reachable and unreachable nodes.

show protocols (IS-IS)

To group a number of protocol show commands according to the specified address family, use the **show protocols** command in EXEC mode.

show protocols [**afi-all** | **ipv4** | **ipv6**] [**all***protocol*]

Syntax Description

afi-all	(Optional) Specifies all address families.
ipv4	(Optional) Specifies an IPv4 address family.
ipv6	(Optional) Specifies an IPv6 address family.
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** • **ospf** • **rip** • **eigrp** For the IPv6 address family, the options are: • **bgp** • **isis** • **ospfv3**

Command Default

If no address family is specified, the default is IPv4.

Command Modes

EXEC

Command History

Release	Modification
Release 3.7.2	This command was introduced.
Release 3.9.0	Support for IPv6 was added

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.

If IPv6 is enabled on an IS-IS instance, the instance is displayed in the **show protocols ipv6** command output. IPv4 IS-IS instances are displayed in the **show protocols ipv4** command output.

When using the **show protocols** command with the **ipv6** or **ipv4** keyword, you get all routing instances in that particular address family—not only IS-IS instances.

Task ID

Task ID	Operations
isis	read
rib	read

Examples

The following example shows the output for the **show protocols** command :

```
Router# show protocols ipv4

IS-IS Router: uut
  System Id: 0000.0000.12a8
  IS Levels: level-1-2
  Manual area address(es):
    49.1515.1515
  Routing for area address(es):
    49.1515.1515
  Non-stop forwarding: Disabled
  Most recent startup mode: Cold Restart
  Topologies supported by IS-IS:
    IPv4 Unicast
      Level-1
        Metric style (generate/accept): Narrow/Narrow
        ISPF Status: Disabled
      Level-2
        Metric style (generate/accept): Narrow/Narrow
        ISPF Status: Disabled
      Redistributing:
        static
      Distance: 115
    IPv6 Unicast
      Level-1
        ISPF Status: Disabled
      Level-2
        ISPF Status: Disabled
      No protocols redistributed
      Distance: 45
  Interfaces supported by IS-IS:
    GigabitEthernet 0/6/0/0 is running actively (active in configuration)
```

This table describes the significant fields shown in the display.

Table 86: show protocols ipv4 Field Descriptions

Field	Description
System ID	Dynamic hostname of the system. The hostname is specified using the **hostname** command. If the dynamic hostname is not known or **hostname dynamic disable** command has been executed, the 6-octet system ID is used.
IS Levels	IS-IS level of the router.
Manual area address(es)	Area addresses configured manually on the originating router.

Field	Description
Routing for area address(es)	Area addresses for which this router provides the routing.
Non-stop forwarding	Status and name of NSF.
Most recent startup mode	Mode in which the most recent startup was performed.
Topologies supported by IS-IS	Address and subaddress family IS-IS are configured.
Metric style	Type, length, and value (TLV) objects accepted by IS-IS. To configure this value, see the metric-style narrow, on page 902, metric-style transition, on page 904, or metric-style wide, on page 906 command.
ISPF status	State of iSPF configuration for this IS-IS instance. Four states exist: • Disabled if iSPF has not been configured but is awaiting a full SPF to compile the topology for use by the iSPF algorithm. • Dormant if iSPF has been configured but is awaiting initial convergence before initializing. • Awake if iSPF has been configured but is awaiting a full SPF to compile the topology for use by the iSPF algorithm. • Active if IS-IS is ready to consider using the iSPF algorithm whenever a new route calculation needs to be run.
Redistributing	IS-IS is configured to redistribute IP static routes into Level 1 or Level 2. The **redistribute** command is used to configure redistribution.
Distance	Administrative distance.
Interfaces supported by IS-IS	Interfaces and their states currently supported by IS-IS. Both operational and configuration status are displayed.

The following example shows how to disable the IPv4 address family, with no output shown for IS-IS IPv4 instances from the **show protocols ipv4** command:

```
Router# configure
Router(config)# router isis uut
Router(config-isis)# no address-family ipv4 unicast
Router(config-isis)# commit

Router# show protocols ipv4
```

Related Commands

Command	Description
metric-style narrow, on page 902	Configures the IS-IS software to generate and accept old-style type, length, and value (TLV) objects.
metric-style transition, on page 904	Configures the IS-IS software to generate and accept both old-style and new-style type length, and value (TLV) objects.
metric-style wide, on page 906	Configures the IS-IS software to generate and accept only new-style type, length, and value (TLV) objects.

Command	Description
redistribute (IS-IS), on page 940	Redistributes routes from one IS-IS instance into another instance.

shutdown (IS-IS)

To disable the Intermediate System-to-Intermediate System (IS-IS) protocol on a particular interface, use the **shutdown** command in interface configuration mode. To re-enable the IS-IS protocol, use the **no** form of this command.

shutdown
no shutdown

Command Default

IS-IS protocol is enabled.

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 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
isis	read, write

Examples

The following example disables the IS-IS protocol on GigabitEthernet interface 0/1/0/1:

```
RP/0/RSP0/CPU0:router(config)# router isis isp
RP/0/RSP0/CPU0:router(config-isis)# interface GigabitEthernet0/1/0/1
RP/0/RSP0/CPU0:router(config-isis-if)# shutdown
```

single-topology

To configure the link topology for IP Version 4 (IPv4) when IP Version 6 (IPv6) is configured, use the **single-topology** command in address family configuration mode. To remove the **single-topology** command from the configuration file and restore the system to its default condition, use the **no** form of this command.

single-topology
no single-topology

Command Default

Performs in multitopology mode in which independent topologies for IPv4 and IPv6 are running in a single area or domain.

Command Modes

IPv6 address family configuration

Command History

Release	Modification
Release 3.9.0	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.

Use the **single-topology** command to allow Intermediate System-to-Intermediate System (IS-IS) for IPv6 to be configured on interfaces along with an IPv4 network protocol. All interfaces must be configured with the identical set of network protocols, and all routers in the IS-IS area (for Level 1 routing) or the domain (for Level 2 routing) must support the identical set of network layer protocols on all interfaces.

When single-topology support for IPv6 is being used, only old-style type, length, and value (TLV) objects may be used and a single shortest path (SPF) individual level is used to compute IPv4 (if configured) and IPv6 routes. The use of a single SPF means that both IPv4 IS-IS and IPv6 IS-IS routing protocols must share a network topology.

To allow link information to be shared between IPv4 and IPv6, you must configure the **single-topology** command for an address family. In single-topology IPv6 mode, the configured metric is always the same for both IPv4 and IPv6.

Task ID

Task ID	Operations
isis	read, write

Examples

The following example shows how to enable single-topology mode for IPv6:

```
RP/0/RSP0/CPU0:router(config)# router isis isp
RP/0/RSP0/CPU0:router(config-isis)# net 49.0000.0000.0001.00
RP/0/RSP0/CPU0:router(config-isis)# address-family ipv6 unicast
RP/0/RSP0/CPU0:router(config-isis-af)# single-topology
```

snmp-server traps isis

snmp-server traps isis {**all** | *traps set*}
no snmp-server traps isis {**all** | *traps set*}

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.

Examples

```
RP/0/RSP0/CPU0:router(config)# snmp-server traps isis

  adjacency-change                isisAdjacencyChange
  all                             Enable all IS-IS traps
  area-mismatch                   isisAreaMismatch
  attempt-to-exceed-max-sequence  isisAttemptToExceedMaxSequence
  authentication-failure          isisAuthenticationFailure
  authentication-type-failure     isisAuthenticationTypeFailure
  corrupted-lsp-detected          isisCorruptedLSPDetected
  database-overload               isisDatabaseOverload
  id-len-mismatch                 isisIDLenMismatch
  lsp-error-detected              isisLSPErrorDetected
  lsp-too-large-to-propagate      isisLSPTooLargeToPropagate
  manual-address-drops            isisManualAddressDrops
  max-area-addresses-mismatch     isisMaxAreaAddressesMismatch
  orig-lsp-buff-size-mismatch     isisOrigLSPBuffSizeMismatch
  own-lsp-purge                   isisOwnLSPPurge
  protocols-supported-mismatch    isisProtocolsSupportedMismatch
  rejected-adjacency              isisRejectedAdjacency
  sequence-number-skip            isisSequenceNumberSkip
  version-skew                    isisVersionSkew

RP/0/RSP0/CPU0:router(config)#snmp-server traps isis all
```

```
RP/0/RSP0/CPU0:router(config)# snmp-server traps isis area-mismatch
lsp-error-detected
```

spf-interval

To customize IS-IS throttling of shortest path first (SPF) calculations, use the **spf-interval** command in address family configuration mode. To restore default values, use the **no** form of this command.

spf-interval [**initial-wait** *initial* | **secondary-wait** *secondary* | **maximum-wait** *maximum*] **. . .** [**level** {**1** | **2**}]
no spf-interval [[**initial-wait** *initial* | **secondary-wait** *secondary* | **maximum-wait** *maximum*] **. . .**] [**level** {**1** | **2**}]

Syntax Description

initial-wait *initial*	Initial SPF calculation delay (in milliseconds) after a topology change. Range is 0 to 120000.
secondary-wait *secondary*	Hold time between the first and second SPF calculations (in milliseconds). Range is 0 to 120000.
maximum-wait *maximum*	Maximum interval (in milliseconds) between two consecutive SPF calculations. Range is 0 to 120000.
level { **1** | **2** }	(Optional) Enables the SPF interval configuration for Level 1 or Level 2 independently.

Command Default

initial-wait *initial* : 50 milliseconds

secondary-wait *secondary* : 200 milliseconds

maximum-wait *maximum* : 5000 milliseconds

Command Modes

Address family 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

SPF calculations are performed only when the topology changes. They are not performed when external routes change.

Use the **spf-interval** command to control how often the software can perform the SPF calculation. The SPF calculation is processor intensive. Therefore, it may be useful to limit how often this calculation is done, especially when the area is large and the topology changes often. Increasing the SPF interval reduces the processor load of the router, but potentially slows the rate of convergence.

Task ID

Task ID	Operations
isis	read, write

Examples

The following example shows how to set the initial SPF calculation delay to 10 milliseconds and the maximum interval between two consecutive SPF calculations to 5000 milliseconds:

```
RP/0/RSP0/CPU0:router(config)# router isis isp
RP/0/RSP0/CPU0:router(config-isis)# address-family ipv4 unicast
RP/0/RSP0/CPU0:router(config-isis-af)# spf-interval initial-wait 10 maximum-wait 5000
```

spf-interval ietf

To set an shortest path first (SPF) interval in IS-IS for SPF calculations, use the **spf-interval ietf** command in the System Admin Config mode. Use the **no** form of this command to enable the fabric bundle port.

spf-interval ietf [**initial-wait** *msec* | **short-wait** *msec* | **long-wait** *msec* | **learn-interval** *msec* | **holddown-interval** *msec*] [**level** { **1** | **2** }]

Syntax Description

spf-interval	Specifies the number of seconds between two consecutive SPF calculations.
ietf	Specifies Internet Engineering Task Force (IETF) RFC standard 8405.
initial-wait *msec*	Initial SPF calculation delay before running a route calculation. The initial-wait must be less than or equal to short-wait. Range is 0 to 120000. The default value is 50 milliseconds.
short-wait *msec*	Short SPF calculation delay before running a route calculation. The short-wait must be less than or equal to long-wait. Range is 0 to 120000. The default value is 200 milliseconds.
long-wait *msec*	Long SPF calculation delay before running a route calculation. Range is 0 to 120000. The default value is 5000 milliseconds.
learn-interval *msec*	Time To Learn interval for running a route calculation. The learn-interval must be less than or equal to holddown-interval. Range is 0 to 120000. The default value is 500 milliseconds.
holddown-interval *msec*	Hold-down interval for running a route calculation. Range is 0 to 120000. The default value is 10000 milliseconds.
level { **1** | **2** }	(Optional) Enables the SPF interval configuration for Level 1 or Level 2 independently.

Command Default

None

Command Modes

System Admin Config mode

Command History

Release	Modification
Release 7.7.1	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.

SPF calculations are performed only when the topology changes. They are not performed when external routes change.

Task ID

Task ID	Operations
is-is	read, write

Example

The following example shows how to configure IETF to set an SPF interval in IS-IS for SPF calculations.

```
Router# configure
Router(config)# router isis isp
Router(config-isis)# address-family ipv4 unicast
Router(config-isis-af)# spf-interval ietf?
initial-wait   Initial delay before running a route calculation [50]
short-wait     Short delay before running a route calculation [200]
long-wait      Long delay before running a route calculation [5000]
learn-interval Time To Learn interval for running a route calculation [500]
holddown-interval  Holddown interval for running a route calculation [10000]
level             Set SPF interval for one level only
Router(config-isis-af)# spf-interval ietf
Router(config-isis-af)#commit
```

The following **show** command displays the output with the new spf-interval algorithm. The output displays the actual delay taken to compute the SPF.

```
Router# show isis ipv4 spf-log last 5 detail
   IS-IS 1 Level 2 IPv4 Unicast Route Calculation Log
                   Time Total Trig.
Timestamp    Type   (ms) Nodes Count First Trigger LSP    Triggers
------------ ----- ----- ----- ----- -------------------- -----------------------
--- Wed Mar 16 2022 ---
15:31:49.763  FSPF     1     6     3         tb5-r4.00-00 LINKBAD PREFIXBAD
  Delay:                 101ms (since first trigger)
                         261177ms (since end of last calculation)
  Trigger Link:          tb5-r2.00
  Trigger Prefix:        34.1.24.0/24
  New LSP Arrivals:      0
  SR uloop:              No
  Next Wait Interval:    200ms
  RIB Batches:           1 (0 critical, 0 high, 0 medium, 1 low)
  Timings (ms):          +--Total--+
                          Real   CPU
    SPT Calculation:         1     1
    Route Update:            0     0
                         ----- -----
```

spf prefix-priority (IS-IS)

To assign a priority to an ISIS prefix for customizing the RIB update sequence, use the**spf prefix-priority** command in address family configuration mode. To restore default values, use the **no** form of this command.

spf prefix-priority [**level** {**1** | **2**}] {**critical** | **high** | **medium**} {*access-list-name* | **tag** *tag*}
no spf prefix-priority [**level** {**1** | **2**}] {**critical** | **high** | **medium**} [*access-list-name* | **tag** *tag*]

Syntax Description

level { **1** | **2** }	(Optional) Enables the assignment of a priority to Level 1 or Level 2 independently.
critical	Assigns a critical priority.
high	Assigns a high priority.
medium	Assigns a medium priority.
access-list-name	Name of an access list.
tag *tag*	Specifies a tag to indicate priority. The *tag* argument range is 1 to 4294967295.

Command Default

By default, IPv4 prefixes with a length of 32 and IPv6 prefixes with a length of 128 are given medium priority. The remaining prefixes are given low priority.

Command Modes

Address family 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **spf prefix-priority** command to change the sequence of prefix updates to the RIB after an SPF is run. ISIS installs prefixes in the RIB according to the following priority order:

Critical > High > Medium > Low

The **spf prefix-priority** command supports prefix lists for the first three priorities. The unmatched prefixes are updated with low priority.

If a **spf prefix-priority** is specified, the default behavior of prioritizing either length 32 or 128 prefixes for IPv4 or IPv6, respectively, as **medium** is disabled.

Task ID

Task ID	Operations
isis	read, write

Examples

The following example shows how to set the prefix priorities:

```
RP/0/RSP0/CPU0:router(config)# ipv4 prefix-list isis-critical-acl
RP/0/RSP0/CPU0:router(config-ipv4_pfx)# 10 permit 0.0.0.0/0 eq 32
!
RP/0/RSP0/CPU0:router(config)# ipv4 prefix-list isis-med-acl
RP/0/RSP0/CPU0:router(config-ipv4_pfx)# 10 permit 0.0.0.0/0 eq 29
!
RP/0/RSP0/CPU0:router(config)# ipv4 prefix-list isis-high-acl
RP/0/RSP0/CPU0:router(config-ipv4_pfx)# 10 permit 0.0.0.0/0 eq 30
!
RP/0/RSP0/CPU0:router(config)# router isis ring
RP/0/RSP0/CPU0:router(config-isis)# address-family ipv4 unicast
RP/0/RSP0/CPU0:router(config-isis-af)# spf prefix-priority critical isis-critical-acl
RP/0/RSP0/CPU0:router(config-isis-af)# spf prefix-priority high isis-high-acl
RP/0/RSP0/CPU0:router(config-isis-af)# spf prefix-priority medium isis-med-acl
```

summary-prefix (IS-IS)

To create aggregate addresses for the Intermediate System-to-Intermediate System (IS-IS) protocol, use the **summary-prefix** command in address family configuration mode. To restore the default behavior, use the **no** form of this command.

Syntax Description

address	Summary address designated for a range of IPv4 addresses. The *address* argument must be in four-part, dotted-decimal notation.
/ *prefix-length*	Length of the IPv4 or IPv6 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.
ipv6-prefix	Summary prefix designated for a range of IPv6 prefixes. The *ipv6-prefix* argument must be in the form documented in RFC 2373, in which the address is specified in hexadecimal using 16-bit values between colons.
level { **1** | **2** }	(Optional) Redistributes routes into Level 1 or Level 2 and summarizes them with the configured address and mask value.
tag *tag*	Sets a tag value. The value range is 1- 4294967295.

Command Default

All redistributed routes are advertised individually.

Both Level 1 and Level 2 are configured if no level is specified.

Command Modes

Address family configuration

Command History

Release	Modification
Release 3.7.2	This command was introduced.
Release 3.9.0	Tag keyword and IPv6 support was added.

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.

Multiple groups of addresses can be summarized for a given level. Routes learned from other routing protocols can also be summarized. The metric used to advertise the summary is the smallest metric of all the more-specific routes. Use the **summary-prefix** command to help reduce the size of the routing table.

This command also reduces the size of the link-state packets (LSPs) and thus the link-state database. It also helps ensure stability, because a summary advertisement depends on many more specific routes. If one more-specific route flaps, in most cases, this flap does not cause a flap of the summary advertisement.

The drawback of summary addresses is that other routes might have less information to calculate the most optimal routing table for all individual destinations.



Note When IS-IS advertises a summary prefix, it automatically inserts the summary prefix into the IP routing table but labels it as a "discard" route entry. Any packet that matches the entry is discarded to prevent routing loops. When IS-IS stops advertising the summary prefix, the routing table entry is removed.

Task ID

Task ID	Operations
isis	read, write

Examples

The following example shows how to redistribute Open Shortest Path First (OSPF) routes into IS-IS:

```
RP/0/RSP0/CPU0:router(config)# router isis isp
RP/0/RSP0/CPU0:router(config-isis)# address-family ipv4 ipv6 unicast
RP/0/RSP0/CPU0:router(config-isis-af)# redistribute ospf 2 level-2
RP/0/RSP0/CPU0:router(config-isis-af)# summary-prefix 10.10.10.10 level-2
RP/0/RSP0/CPU0:router(config-isis-af)# summary-prefix 10.10.10.10
```

suppressed

To allow an IS-IS interface to participate in forming adjacencies without advertising connected prefixes in the system link-state packets (LSPs), use the **suppressed** command in interface configuration mode. To enable advertising connected prefixes, use the **no** form of this command.

suppressed
no suppressed

Command Default

Interface is active.

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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **suppressed** command to reduce the number of routes that IS-IS has to maintain, improving convergence times after an isolated failure. Improvement is noticeable if the command is used widely throughout the network. Other routers in the domain do not install routes to the affected connected prefixes.

Task ID

Task ID	Operations
isis	read, write

Examples

The following example shows how to disable the advertisement of connected prefixes on GigabitEthernet interface 0/1/0/1:

```
RP/0/RSP0/CPU0:router(config)# router isis isp
RP/0/RSP0/CPU0:router(config-isis)# interface GigabitEthernet /1/0/1
RP/0/RSP0/CPU0:router(config-isis-if)# suppressed
```

Related Commands

Command	Description
passive (IS-IS), on page 932	Suppresses S-IS packets on an interface.

tag (IS-IS)

To associate and advertise a tag with the prefix of an IS-IS interface, use the **tag** command in interface address family configuration mode. To restore the default behavior, use the **no** form of this command.

tag *tag*
no tag [*tag*]

Syntax Description

tag	Interface tag. Range is 1 to 4294967295.

Command Default

Default is that no tag is associated and advertised.

Command Modes

Interface address family 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 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
isis	read, write

Examples

The following example shows how to associate and advertise an interface tag:

```
RP/0/RSP0/CPU0:router(config)# router isis isp
RP/0/RSP0/CPU0:router(config-isis)# interface GigabitEthernet 0/3/0/0
RP/0/RSP0/CPU0:router(config-isis-if)# address-family ipv4 unicast
RP/0/RSP0/CPU0:router(config-isis-if-af)# tag 234
```

Related Commands

Command	Description
spf prefix-priority (IS-IS), on page 1026	Assigns a priority to an ISIS prefix for customizing the RIB update sequence.

topology-id

To differentiate one topology in the domain from another while configuring a multicast routing table, use the **topology-id** command in Intermediate System-to-Intermediate System (IS-IS) address family configuration submode. To disable the topology use the **no** form of the command.

topology-id *isis-multicast-topology-id-number*
no topology-id *isis-multicast-topology-id-number*

Syntax Description

isis-multicast-topology-id-number	ID number for a specific IS-IS multicast topology. Range is 6 to 4095.

Command Default

No topology is associated with a routing table by default.

Command Modes

IS-IS address family 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 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
isis	read, write

Examples

The following example shows how to differentiate a topology from another in the multicast routing table in IS-IS routing:

```
RP/0/RSP0/CPU0:router(config)# router  isis isp
RP/0/RSP0/CPU0:router(config-isis)# address-family ipv4 multicast topology   green
RP/0/RSP0/CPU0:router(config-isis-af)# topology-id 2666
```

Related Commands

Command	Description
address-family multicast topology (IS-IS), on page 805	Used in conjunction with the **topology-id** command, enables a multicast topology globally when configuring Intermediate System-to-Intermediate System (IS-IS) routing.

trace (IS-IS)

To set the the IS-IS buffer size, use the **trace** command in router configuration mode. To return to the default value, use the **no** form of this command.

trace [**detailed** | **severe** | **standard**] *max-trace-entries*
no trace [**detailed** | **severe** | **standard**]

Syntax Description

detailed	Specifies the buffer size for detailed traces. Range is
severe	Specifies the buffer size for severe traces. Range is
standard	Specifies the buffer size for standard traces. Range is
max-trace-entries	Sets the maximum number of trace entries. Range is 1-20000

Command Default

None

Command Modes

Router IS-IS configuration

Command History

Release	Modification
Release 3.9.0	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
isis	read, write

Examples

The following example shows how to set the isis buffer size for severe traces to 1200:

```
RP/0/RSP0/CPU0:router(config)#router isis isp
RP/0/RSP0/CPU0:router(config-isis)#trace sever 1200
```

ucmp (IS-IS)

To enable unequal cost multipath (UCMP) calculation for Intermediate System-to-Intermediate System (IS-IS), use the **ucmp** command in IS-IS address family configuration mode. To disable UCMP for IS-IS, use the **no** form of this command.

ucmp [**prefix-list** *prefix-list-name*] [**variance** *value*]
no ucmp

Syntax Description

prefix-list	(Optional) Specifies prefix-list name to filter UCMP paths based on prefixes.
prefix-list-name	Name of the prefix-list to be specified to filter UCMP paths.
variance	(Optional) Specifies variance parameter to filter UCMP paths based on cost.
variance-value	Variance value. The variance value is expressed in terms of percentage of the Primary path metric. Range is 101 to 10000. Default variance value is 200.

Command Default

UCMP is disabled.

Command Modes

IS-IS IPv4 address family

IS-IS IPv6 address family

Command History

Release	Modification
Release 4.3.0	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.

Enabling the **ucmp** command will make the router calculate UCMP nexthops for all the prefixes in the routing table.

The UCMP path calculation can be controlled such that UCMP nexthops are calculated only for a certain set of prefixes. Use the **ucmp** command with the **prefix-set** option to enable this functionality.

If there are multiple UCMP nexthops with various metrics, then the selection of the number of UCMP nexthops is controlled by the **variance** option in the UCMP command. The variance value is expressed in terms of percentage of the primary path metric. For example, if variance value is 150 and primary path metric is 100, then select all the UCMP nexthops with metrics from 101 to 150.

Task ID

Task ID	Operation
isis	read, write

This example shows how to enable calculation of UCMP nexthops for all the prefixes in the routing table:

```
Router#configure
Router(config)#router isis isp
Router(config-isis)#address-family ipv6 unicast
Router(config-isis-af)#ucmp
```

This example shows how to enable calculation of UCMP nexthops for a set of prefixes in the prefix-list *list1*:

```
Router#configure
Router(config)#router isis isp
Router(config-isis)#address-family ipv6 unicast
Router(config-isis-af)#ucmp prefix-list list1
```

This example shows how to enable calcualtion of UCMP mexthops with variance value *120*:

```
Router#configure
Router(config)#router isis isp
Router(config-isis)#address-family ipv6 unicast
Router(config-isis-af)#ucmp variance 120
```

Related Commands

Command	Description
ucmp delay-interval	Specifies delay between primary SPF completion and start of UCMP computation.
ucmp exclude interface	Excludes an interface from unequal cost multipath (UCMP) computation.
bandwidth	Configures the bandwidth of an interface.

ucmp delay-interval

To specify delay between primary SPF completion and start of UCMP computation, use the **ucmp delay-interval** command in IS-IS address family configuration mode. To disable this functionality, use the **no** form of this command.

ucmp delay-interval *delay-interval*
no ucmp delay-interval

Syntax Description

delay-interval	Delay interval value in milliseconds. Range is 100 to 65535 milliseconds. The default value for the interval is 100 milliseconds.

Command Default

UCMP delay interval is set to 100 milliseconds.

Command Modes

IS-IS IPv4 address family

IS-IS IPv6 address family

Command History

Release	Modification
Release 4.3.0	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
isis	read, write

This example shows how to set the delay between primary SPF completion and the start of UCMP computation, to 800 milliseconds:

```
Router#configure
Router(config)#router isis isp
Router(config-isis)#address-family ipv6 unicast
Router(config-isis-af)#ucmp delay-interval 800
```

Related Commands

Command	Description
ucmp is-is	Enables unequal cost multipath (UCMP) calculation for Intermediate System-to-Intermediate System (IS-IS).

Command	Description
ucmp exclude interface	Excludes an interface from unequal cost multipath (UCMP) computation.
bandwidth	Configures the bandwidth of an interface.

ucmp exclude interface

To exclude an interface from unequal cost multipath (UCMP) computation, use the **ucmp exclude interface** command in IS-IS address family configuration mode. To disable this functionality, use the **no** form of this command.

ucmp exclude interface *type interface-path-id*
no ucmp exclude interface *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

None

Command Modes

IS-IS IPv4 address family

IS-IS IPv6 address family

Command History

Release	Modification
Release 4.3.0	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.

Use the ucmp exclude interface command to exclude an interface from being selected as a UCMP backup path.

Task ID

Task ID	Operation
isis	read, write

This example shows how to exclude interface POS 0/3/0/1 from UCMP computation:

```
Router#configure
Router(config)#router isis isp
Router(config-isis)#address-family ipv6 unicast
Router(config-isis-af)#ucmp exclude interface POS 0/3/0/1
```

Related Commands

Command	Description
ucmp is-is	Enables unequal cost multipath (UCMP) calculation for Intermediate System-to-Intermediate System (IS-IS).
ucmp delay-interval	Specifies delay between primary SPF completion and start of UCMP computation.
bandwidth	Configures the bandwidth of an interface.



OSPF Commands

This module describes the commands used to configure and monitor the Open Shortest Path First (OSPF) routing protocol.

For detailed information about OSPF concepts, configuration tasks, and examples, see the *Implementing OSPF on* Cisco ASR 9000 Series Router module in the *Routing Configuration Guide for Cisco ASR 9000 Series Routers* .

- address-family (OSPF), on page 1044
- adjacency stagger (OSPF), on page 1045
- area (OSPF), on page 1047
- authentication (OSPF), on page 1049
- authentication-key (OSPF), on page 1051
- auto-cost (OSPF), on page 1053
- capability opaque disable, on page 1055
- capability vrf-lite (OSPFv2), on page 1056
- clear ospf process, on page 1057
- clear ospf redistribution, on page 1059
- clear ospf routes, on page 1061
- clear ospf statistics, on page 1062
- cost (OSPF), on page 1064
- cost-fallback (OSPF), on page 1066
- database-filter all out (OSPF), on page 1068
- dead-interval (OSPF), on page 1069
- default-cost (OSPF), on page 1071
- default-information originate (OSPF), on page 1073
- default-metric (OSPF), on page 1075
- demand-circuit (OSPF), on page 1077
- disable-dn-bit-check, on page 1079
- distance (OSPF), on page 1080
- distance ospf, on page 1083
- distribute-list, on page 1085
- domain-id (OSPF), on page 1087
- domain-tag, on page 1089
- exchange-timer, on page 1090
- fast-reroute (OSPFv2), on page 1092

- fast-reroute per-link exclude interface, on page 1094
- fast-reroute per-prefix exclude interface (OSPFv2), on page 1096
- fast-reroute per-prefix lfa-candidate (OSPFv2), on page 1098
- fast-reroute per-prefix remote-lfa (OSPFv2), on page 1099
- fast-reroute per-prefix ti-lfa, on page 1101
- fast-reroute per-prefix use-candidate-only (OSPFv2), on page 1102
- flood-reduction (OSPF), on page 1103
- hello-interval (OSPF), on page 1105
- ignore lsa mospf, on page 1107
- interface (OSPF), on page 1108
- log adjacency changes (OSPF), on page 1110
- loopback stub-network , on page 1111
- lpts punt excessive-flow-trap penalty-timeout ospf, on page 1112
- max external-lsa, on page 1113
- max-lsa, on page 1115
- max-metric, on page 1118
- maximum interfaces (OSPF), on page 1121
- maximum paths (OSPF), on page 1123
- maximum redistributed-prefixes (OSPF), on page 1125
- message-digest-key, on page 1127
- microloop avoidance (OSPF), on page 1130
- mpls ldp auto-config (OSPF), on page 1131
- mpls ldp sync (OSPF), on page 1132
- mpls traffic-eng (OSPF), on page 1133
- mpls traffic-eng igp-intact (OSPF), on page 1135
- mpls traffic-eng multicast-intact (OSPF), on page 1137
- mpls traffic-eng ldp-sync-update (OSPF), on page 1138
- mpls traffic-eng router-id (OSPF), on page 1139
- mtu-ignore (OSPF), on page 1141
- multi-area-interface, on page 1143
- neighbor (OSPF), on page 1145
- neighbor database-filter all out, on page 1147
- network (OSPF), on page 1148
- nsf (OSPF), on page 1150
- nsf flush-delay-time (OSPF), on page 1152
- nsf interval (OSPF), on page 1153
- nsf lifetime (OSPF), on page 1154
- nsr (OSPF), on page 1155
- nssa (OSPF), on page 1157
- ospf name-lookup, on page 1159
- packet-size (OSPF), on page 1160
- passive (OSPF), on page 1162
- prefix-suppression, on page 1164
- priority (OSPF), on page 1166
- protocol shutdown, on page 1168
- queue dispatch flush-lsa, on page 1169

- queue dispatch incoming, on page 1171
- queue dispatch rate-limited-lsa, on page 1173
- queue dispatch spf-lsa-limit, on page 1175
- queue limit, on page 1176
- range (OSPF), on page 1178
- redistribute (OSPF), on page 1180
- retransmit-interval (OSPF), on page 1185
- route-policy (OSPF), on page 1187
- router-id (OSPF), on page 1188
- router ospf (OSPF), on page 1190
- security ttl (OSPF), on page 1192
- segment-routing sr-prefer prefix-list, on page 1194
- sham-link, on page 1196
- show lpts punt excessive-flow-trap ospf, on page 1198
- show ospf, on page 1199
- show ospf border-routers, on page 1202
- show ospf database, on page 1204
- show ospf flood-list, on page 1219
- show ospf interface, on page 1221
- show ospf mpls traffic-eng, on page 1224
- show ospf message-queue, on page 1229
- show ospf neighbor, on page 1232
- show ospf request-list, on page 1241
- show ospf retransmission-list, on page 1244
- show ospf routes, on page 1246
- show ospf sham-links, on page 1251
- show ospf summary-prefix, on page 1254
- show ospf virtual-links, on page 1256
- show protocols (OSPF), on page 1258
- snmp context (OSPF), on page 1261
- snmp trap (OSPF), on page 1263
- snmp trap rate-limit (OSPF), on page 1264
- spf prefix-priority (OSPFv2), on page 1265
- stub (OSPF), on page 1267
- summary-prefix (OSPF), on page 1269
- timers lsa group-pacing, on page 1271
- timers lsa min-arrival, on page 1272
- timers throttle lsa all (OSPF), on page 1273
- timers throttle spf (OSPF), on page 1276
- transmit-delay (OSPF), on page 1278
- virtual-link (OSPF), on page 1280
- vrf (OSPF), on page 1282

address-family (OSPF)

To enter address family configuration mode for Open Shortest Path First (OSPF), use the **address-family** command in the appropriate mode. To disable address family configuration mode, use the **no** form of this command.

address-family ipv4 [**unicast**]
no address-family ipv4 [**unicast**]

Syntax Description

ipv4	Specifies IP Version 4 (IPv4) address prefixes.
unicast	(Optional) Specifies unicast address prefixes.

Command Default

An address family is not specified.

Command Modes

Router configuration

VRF 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

OSPF version 2 automatically provides routing services for IPv4 unicast topologies, so this command is redundant.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to configure the OSPF router process with IPv4 unicast address prefixes:

```
RP/0/RSP0/CPU0:router(config)# router ospf 1
RP/0/RSP0/CPU0:router(config-ospf)# address-family ipv4 unicast
```

adjacency stagger (OSPF)

To configure staggering of OSPF adjacency during reload, process restart, and process clear, use the **adjacency stagger** command in router configuration mode. To turn off adjacency staggering, either use the **disable** keyword or use the **no** form of this command.

adjacency stagger {**disable** | *initial-num-nbr* *max-num-nbr*}
no adjacency stagger

Syntax Description

disable	Disables adjacency staggering.
initial-num-nbr	The initial number of simultaneous neighbors allowed to form adjacency to FULL in any area to bring up to FULL after a router reload, OSPF process restart, or OSPF process clear. Range is 1-65535. Default is 2.
max-num-nbr	The subsequent number of simultaneous neighbors allowed to form adjacency, per OSPF instance, after the initial set of OSPF neighbors have become FULL. Range is 1-65535. Default is 64.

Command Default

OSPF adjacency staggering is enabled.

Command Modes

Router configuration

Command History

Release	Modification
Release 3.9.0	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.

Staggering of the OSPF adjacency during reload, process restart (without NSR or graceful-restart), and process clear reduces the overall adjacency convergence time.

Initially, allow 2 (configurable) neighbors to form adjacency to FULL per area. After the first adjacency reaches FULL, up to 64 (configurable) neighbors can form adjacency simultaneously for the OSPF instance (all areas). However, areas without any FULL adjacency is restricted by the initial area limit.

Note Adjacency stagger and OSPF nonstop forwarding (NSF) are mutually exclusive. Adjacency stagger will not be activated if **nsf** is configured under router ospf configuration.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to configure adjacency stagger for a 2 neighbors initially and for a maximum of 3 neighbors:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router ospf 1
RP/0/RSP0/CPU0:router(config-ospf)# adjacency stagger 2 3
```

area (OSPF)

To configure an Open Shortest Path First (OSPF) area, use the **area** command in the appropriate mode. To terminate an OSPF area, use the **no** form of this command.

area *area-id*
no area *area-id*

Syntax Description

area-id	Identifier of an OSPF area. The *area-id* argument can be specified as either a decimal value or an IP address (dotted decimal) format. Range is 0 to 4294967295.

Command Default

No OSPF area is defined.

Command Modes

Router configuration

VRF 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **area** command to explicitly configure an area. Commands configured under the area configuration mode (such as the **interface** [OSPF] and **authentication** commands), are automatically bound to that area.

To modify or remove the area, the *area-id* argument format must be the same as the format used when creating the area. Otherwise, even if the actual 32-bit value matches, the area is not matched. For example, if you create an area with an *area-id* of 10 it would not match an *area-id* of 0.0.0.10.

Note To remove the specified area from the router configuration, use the **no area** *area-id* command. The **no area** *area-id* command removes the area and all area options, such as **authentication** , **default-cost** , **nssa** , **range** , **stub** , **virtual-link** , and **interface.**

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to configure area 0 and GigabitEthernet interface 0/2/0/0. GigabitEthernet interface 0/2/0/0 is bound to area 0 automatically.

```
RP/0/RSP0/CPU0:router# configure
```

```
RP/0/RSP0/CPU0:router(config)# router ospf 1
RP/0/RSP0/CPU0:router(config-ospf)# area 0
RP/0/RSP0/CPU0:router(config-ospf-ar)# interface GigabitEthernet 0/2/0/0
```

authentication (OSPF)

To enable plain-text authentication, Message Digest 5 (MD5) authentication, Keychain-based authentication or null authentication for an Open Shortest Path First (OSPF) interface, use the **authentication** command in the appropriate mode. To remove such authentication, use the **no** form of this command.

authentication [**message-digest** [**keychain** *keychain*] | **null**]
no authentication

Syntax Description

message-digest	(Optional) Specifies that MD5 is used.
keychain *keychain*	(Optional) Specifies a keychain name.
null	(Optional) Specifies that no authentication is used. Useful for overriding password or MD5 authentication if configured for an area.

Command Default

If this command is not specified in interface configuration mode, then the interface adopts the authentication parameter specified by the area.

If this command is not specified in area configuration mode, then the interface adopts the authentication parameter specified for the process.

If this command is not specified at any level, then the interface does not use authentication.

If no keyword is specified, plain text authentication is used.

Command Modes

Interface configuration

Area configuration

Router configuration

Virtual-link configuration

VRF configuration

Multi-area interface configuration

Sham-link 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **authentication** command to specify an authentication type for the interface, which overrides the authentication specified for the area to which this interface belongs. If this command is not included in the configuration file, the authentication configured in the area to which the interface belongs is assumed (as specified by the area **authentication** command).

The authentication type and password must be the same for all OSPF interfaces that are to communicate with each other through OSPF. If you specified plain text authentication, use the **authentication-key** command to specify the plain text password.

If you enable MD5 authentication with the **message-digest** keyword, you must configure a key with the **message-digest-key** interface command.

To manage the rollover of keys and enhance MD5 authentication for OSPF, you can configure a container of keys called a keychain with each key comprising the following attributes: generate/accept time, key identification, and authentication algorithm. The keychain management feature is always enabled.



Note Changes to the system clock will impact the validity of the keys in the existing configuration.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to set authentication for areas 0 and 1 of OSPF routing process 201. Authentication keys are also provided.

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router ospf 201
RP/0/RSP0/CPU0:router(config-ospf)# router-id 10.1.1.1
RP/0/RSP0/CPU0:router(config-ospf)# area 0
RP/0/RSP0/CPU0:router(config-ospf-ar)# authentication
RP/0/RSP0/CPU0:router(config-ospf-ar)# interface GigabitEthernet 0/1/0/1
RP/0/RSP0/CPU0:router(config-ospf-ar-if)# authentication-key mykey
RP/0/RSP0/CPU0:router(config-ospf-ar-if)# exit
RP/0/RSP0/CPU0:router(config-ospf)# area 1
RP/0/RSP0/CPU0:router(config-ospf-ar)# authentication
RP/0/RSP0/CPU0:router(config-ospf-ar)# interface GigabitEthernet 0/1/0/0
RP/0/RSP0/CPU0:router(config-ospf-ar-if)# authentication-key mykey1
```

The following example shows how to configure use of an authentication keychain:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router ospf 201
RP/0/RSP0/CPU0:router(config-ospf)# router-id 10.1.1.1
RP/0/RSP0/CPU0:router(config-ospf)# authentication message-digest keychain mykeychain
```

Related Commands

Command	Description
authentication-key (OSPF), on page 1051	Assigns a password to be used by neighboring routers that are using the simple password authentication of OSPF.
message-digest-key, on page 1127	Specifies a key used with OSPF MD5 authentication.

authentication-key (OSPF)

To assign a password to be used by neighboring routers that are using the Open Shortest Path First (OSPF) simple password authentication, use the **authentication-key** command in the appropriate mode. To remove a previously assigned OSPF password, use the **no** form of this command.

authentication-key [**clear** | **encrypted**] *password*
no authentication-key

Syntax Description

clear	(Optional) Specifies that the key be clear text.
encrypted	(Optional) Specifies that the key be encrypted using a two-way algorithm.
password	Any contiguous string up to 8 characters in length that can be entered from the keyboard. For example, *mypswd2*.

Command Default

If this command is not specified in interface configuration mode, then the interface adopts the OSPF password parameter specified by the area.

If this command is not specified in area configuration mode, then the interface adopts the OSPF password parameter specified for the process.

If this command is not specified at any level, then no password is specified.

Clear is the default if the **clear** or **encrypted** keyword is not specified.

Command Modes

Interface configuration

Area configuration

Router configuration

Virtual-link configuration

VRF configuration

Multi-area configuration

Sham-link 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

The password created by this command is inserted directly into the OSPF header when the Cisco IOS XR software originates routing protocol packets. A separate password can be assigned to each network on an individual interface basis. All neighboring routers on the same network must have the same password to be able to exchange OSPF information.

The **authentication-key** command must be used with the **authentication** command. If the **authentication** command is not configured, the password provided by the **authentication-key** command is ignored and no authentication is adopted by the OSPF interface.



Note The **authentication-key** command cannot be used with the **authentication** command when the **message-digest** or **null** keyword is configured.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to configure an authentication password as the string yourpass:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router ospf 201
RP/0/RSP0/CPU0:router(config-ospf)# authentication-key yourpass
```

Related Commands

Command	Description
authentication (OSPF), on page 1049	Specifies authentication type.

auto-cost (OSPF)

To control how the Open Shortest Path First (OSPF) protocol calculates default metrics for the interface, use the **auto-cost** command in the appropriate mode. To revert to the default reference bandwidth, use the **no** form of this command.

auto-cost {**reference-bandwidth** *mbps* | **disable**}
no auto-cost {**reference-bandwidth** | **disable**}

Syntax Description

reference-bandwidth *mbps*	Specifies a rate in Mbps (bandwidth). Range is 1 to 4294967.
disable	Assigns a cost based on interface type.

Command Default

mbps : 100 Mbps

Command Modes

Router configuration

VRF 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

By default OSPF calculates the OSPF metric for an interface according to the bandwidth of the interface.

The OSPF metric is calculated as the *mbps* value divided by bandwidth, with *mbps* equal to 108 by default.

If you have multiple links with high bandwidth (such as OC-192), you might want to use a larger number to differentiate the cost on those links. That is, the metric calculated using the default *mbps* value is the same for all high-bandwidth links.

Recommended usage of cost configuration for OSPF interfaces with high bandwidth is to be consistent: Either explicitly configure (by using the **cost** command) or choose the default (by using the **auto-cost** command).

The value set by the **cost** command overrides the cost resulting from the **auto-cost** command.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to set the reference value for the auto cost calculation to 1000 Mbps:

```
RP/0/RSP0/CPU0:router# configure
```

```
RP/0/RSP0/CPU0:router(config)# router ospf 1
RP/0/RSP0/CPU0:router(config-ospf)# auto-cost reference-bandwidth 1000
```

Related Commands

Command	Description
cost (OSPF), on page 1064	Explicitly specifies the cost of the interface (network) for OSPF path calculation.

capability opaque disable

To prevent Multiprotocol Label Switching traffic engineering (MPLS TE) topology information flooded to the network through opaque LSAs, use the **capability opaque disable** command in the appropriate mode. To restore MPLS TE topology information flooded through opaque LSAs to the network, use the **no** form of the command.

capability opaque disable
no capability opaque disable

Command Default

Opaque LSAs are allowed.

Command Modes

Router configuration

VRF 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

The **capability opaque disable** command prevents flooded MPLS TE information (Types 1 and 4) through opaque LSAs of all scope (Types 9, 10, and 11).

Control opaque LSA support capability must be enabled for OSPF to support MPLS TE.

The MPLS TE topology information is flooded to the area through opaque LSAs by default.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to prevent OSPF from supporting opaque services:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router ospf 1
RP/0/RSP0/CPU0:router(config-ospf)# capability opaque disable
```

capability vrf-lite (OSPFv2)

To ignore DN bit in LSAs received from peers in the given VRF and to disable automatic ABR status in that VRF, use the **capability vrf-lite** command in OSPFv2 VRF configuration mode. To disable ignoring the DN bit in LSAs and to re-enable automatic ABR status in the VRF, use the **no** form of this command.

capability vrf-lite
no capability vrf-lite

Syntax Description

This command has no keywords or arguments.

Command Default

Disabled

Command Modes

OSPFv2 VRF configuration

Command History

Release	Modification
Release 4.3.0	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.

Use the **capability vrf-lite** command when CE routers (sometimes called multi-vrf CE routers) are directly connected through interfaces associated with a VRF, but without being connected to other PEs through the MPLS/VPN BGP Backbone. Configure the **capability vrf-lite** command for CE routers.

When OSPFv2 is enabled in a VRF, the router is always an ABR. With the **capability vrf-lite** command is enabled, the router becomes an ABR only if it is connected to area 0 (backbone area), and there are other (non-backbone) areas enabled on this router in the given VRF.

Note Routes may be re-introduced to the VPN backbone when this command is used.

Task ID

Task ID	Operation
ospf	read, write

This example shows how to enable vrf-lite capability for OSPF instance *100* under *vrf1*

```
Router#configure
Router(config)#router ospf 100
Router(config-ospf)#vrf vrf1
Router(config-ospf-vrf)#capability vrf-lite
```

clear ospf process

To reset an Open Shortest Path First (OSPF) router process without stopping and restarting it, use the **clear ospf process** command in EXEC configuration mode.

clear ospf [*process-name* [**vrf** {*vrf-name* | **all**}]] **process**

Syntax Description

process-name	(Optional) Name that uniquely identifies an OSPF routing process. The process name is defined by the **router ospf** command. If this argument is included, only the specified routing process is affected. Otherwise, all OSPF processes are reset.
vrf	(Optional) An OSPF VPN routing and forwarding (VRF) instance.
vrf-name	(Optional) Name of the OSPF VRF instance to be reset.
all	(Optional) Resets all OSPF VRF instances.

Command Default

No default behavior or value

Command Modes

EXEC 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

When the OSPF router process is reset, OSPF releases all resources allocated, cleans up the internal database, and shuts down and restarts all interfaces that belong to the process.

Note The **clear ospf process** command may change the router ID unless the OSPF router ID is explicitly configured through the router-id (OSPF), on page 1188 command.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to reset all OSPF processes:

```
RP/0/RSP0/CPU0:router# clear ospf process
```

The following example shows how to reset the OSPF 1 process:

```
RP/0/RSP0/CPU0:router# clear ospf 1 process
```

Related Commands

Command	Description
router ospf (OSPF), on page 1190	Configures an OSPF routing process.
router-id (OSPF), on page 1188	Configures a router ID for the OSPF process.

clear ospf redistribution

To clear all routes redistributed from other protocols out of the Open Shortest Path First (OSPF) routing table, use the **clear ospf redistribution** command in EXEC configuration mode.

clear ospf [*process-name* [**vrf** {*vrf-name* | **all**}]] **redistribution**

Syntax Description

process-name	(Optional) Name that uniquely identifies an OSPF routing process. The process name is defined by the **router ospf** command. If this argument is included, only the specified routing process is affected. Otherwise, all OSPF routes are cleared.
vrf	(Optional) OSPF VPN routing and forwarding (VRF) instance.
vrf-name	(Optional) Name of the OSPF VRF instance to be reset.
all	(Optional) Resets all OSPF VRF instances.

Command Default

No default behavior or value

Command Modes

EXEC 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **clear ospf redistribution** command to cause the routing table to be read again. OSPF regenerates and sends Type 5 and Type 7 link-state advertisements (LSAs) to its neighbors. If an unexpected route has appeared in the OSPF redistribution, using this command corrects the issue.

Note Use of this command can cause a significant number of LSAs to flood the network. We recommend that you use this command with caution.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to clear all redistributed routes across all processes from other protocols:

```
RP/0/RSP0/CPU0:router# clear ospf redistribution
```

clear ospf routes

To clear all Open Shortest Path First (OSPF) routes from the OSPF routing table, use the **clear ospf routes** command in EXEC configuration mode.

clear ospf [*process-name* [**vrf** {*vrf-name* | **all**}]] **routes**

Syntax Description

process-name	(Optional) Name that uniquely identifies an OSPF routing process. The process name is defined by the **router ospf** command. If this argument is included, only the specified routing process is affected. Otherwise, all OSPF routes are cleared.
vrf	(Optional) OSPF VPN routing and forwarding (VRF) instance.
vrf-name	(Optional) Name of the OSPF VRF instance to be reset.
all	(Optional) Resets all OSPF VRF instances.

Command Default

No default behavior or value

Command Modes

EXEC 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 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
ospf	read, write

Examples

The following example shows how to clear all OSPF routes from the OSPF routing table and recompute valid routes. When the OSPF routing table is cleared, OSPF routes in the global routing table are also recalculated.

```
RP/0/RSP0/CPU0:router# clear ospf routes
```

Related Commands

Command	Description
router ospf (OSPF), on page 1190	Configures an OSPF routing process.

clear ospf statistics

To clear the Open Shortest Path First (OSPF) statistics of neighbor state transitions, use the **clear ospf statistics** command in EXEC configuration mode.

clear ospf [*process-name* [**vrf** {*vrf-name* | **all**}]] **statistics** [**neighbor** [*type interface-path-id*] [*ip-address*]]

Syntax Description

process-name	(Optional) Name that uniquely identifies an OSPF routing process. The process name is defined by the **router ospf** command. If this argument is included, only the specified routing process is affected. Otherwise, all OSPF statistics of neighbor state transitions are cleared.
vrf	(Optional) OSPF VPN routing and forwarding (VRF) instance.
vrf-name	(Optional) Name of the OSPF VRF instance to be reset.
all	(Optional) Resets all OSPF VRF instances.
neighbor	(Optional) Clears the state transition counters of the specified neighbor only.
type	(Optional) Interface type. For more information, use the question mark (?) online help function.
interface-path-id	(Optional) Physical interface or virtual interface. 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.
ip-address	(Optional) IP address of a specified neighbor for whom you want to clear the state transition counter.

Command Default

No default behavior or value

Command Modes

EXEC 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **clear ospf statistics** command to reset OSPF counters. Reset is useful to detect changes in counter values.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to reset the OSPF transition state counters for all neighbors on Packet-over-SONET/SDH (POS) interface 0/2/0/0:

```
RP/0/RSP0/CPU0:router# clear ospf statistics neighbor POS 0/2/0/0
```

Related Commands

Command	Description
router ospf (OSPF), on page 1190	Configures an OSPF routing process.

cost (OSPF)

To explicitly specify the interface (network) for Open Shortest Path First (OSPF) path calculation, use the **cost** command in the appropriate mode. To remove the cost, use the **no** form of this command.

cost *cost*
no cost

Syntax Description

cost	Unsigned integer value expressed as the link-state metric. Range is 1 to 65535.

Command Default

If this command is not specified in interface configuration mode, then the interface adopts the cost parameter specified by the area.

If this command is not specified in area configuration mode, then the interface adopts the cost parameter specified for the process.

If this command is not specified at any level, then the cost is calculated by the **auto-cost** command.

Command Modes

Interface configuration

Area configuration

Router configuration

VRF configuration

Multi-area configuration

Sham-link 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

The link-state metric is advertised as the link cost in the router link advertisement. Cisco IOS XR software does not support type of service (ToS), so you can assign only one cost for each interface.

In general, the path cost is calculated using the following formula:

108 / bandwidth (the default auto cost is set to 100 Mbps)

This calculation is the default reference bandwidth used by the auto-costing calculation which establishes the interface auto-cost The **auto-cost** command can set this reference bandwidth to some other value. The **cost** command is used to override the auto-costing calculated default value for interfaces.

Using this formula, the default path cost is 1 for any interface that has a link bandwidth of 100 Mbps or higher. If this value does not suit the network, configure the reference bandwidth for auto calculating costs based on the link bandwidth.

The value set by the **cost** command overrides the cost resulting from the **auto-cost (OSPF)** command.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to set the cost value to 65 for GigabitEthernet interface 0/1/0/1:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router ospf 1
RP/0/RSP0/CPU0:router(config-ospf)# area 0
RP/0/RSP0/CPU0:router(config-ospf-ar)# interface GigabitEthernet 0/1/0/1
RP/0/RSP0/CPU0:router(config-ospf-ar-if)# cost 65
```

Related Commands

Command	Description
auto-cost (OSPF), on page 1053	Controls how the OSPF protocol calculates default metrics for the interface.

cost-fallback (OSPF)

To apply higher cost than the normal interface cost when the cumulative bandwidth of a bundle interface goes below the threshold specified and to revert to the original cost if the cumulative bandwidth goes above the configured threshold, use the **cost-fallback** command. To remove the cost-fallback, use the **no** form of this command.

cost-fallback cost threshold bandwidth
no cost-fallback

Syntax Description

cost **threshold**	Unsigned integer value expressed as the link-state metric. Range is 1 to 65535, but typically, cost-fallback value is supposed to be set to a value higher than the normal cost.
bandwidth	Unsigned integer value expressed in Mbits per second. Range is 1 to 4294967.

Command Default

If this command is not specified in interface configuration mode, the currently effective interface cost takes effect even when the cumulative bandwidth goes down below the maximum bandwidth. Unlike the interface cost command, this cost-fallback command is available only under interface configuration mode; it is not available in area or process level. Unlike other interface specific parameters, no inheritance will take place from area or process level if this command is not specified at interface level.

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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

The fallback cost must be set to a higher value than the normal interface cost. The motivation of setting the fallback cost is to cost out an interface or disfavor an interface without shutting it down when its cumulative bandwidth goes below the user specified threshold, so that the traffic can take an alternative path. The normal interface cost will take over when the cumulative bandwidth reaches or exceeds user-specified threshold.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to set the cost-fallback value:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)#router ospf 100
RP/0/RSP0/CPU0:router(config-ospf)#router-id 2.2.2.2
RP/0/RSP0/CPU0:router(config-ospf)#area 0
```

```
RP/0/RSP0/CPU0:router(config-ospf-ar)#interface bundle-a
RP/0/RSP0/CPU0:router(config-ospf-ar-if)#cost-fallback 1000 threshold 300
```

Related Commands

Command	Description
auto-cost (OSPF), on page 1053	Controls how the OSPF protocol calculates default metrics for the interface.
cost (OSPF), on page 1064	Specifies the cost of the interface (network) for OSPF path calculation.

database-filter all out (OSPF)

To filter outgoing link-state advertisements (LSAs) to an Open Shortest Path First (OSPF) interface, use the **database-filter all out** command in the appropriate mode. To restore the forwarding of LSAs to the interface, use the **disable** form of the command.

database-filter all out [**disable** | **enable**]

Syntax Description

disable	(Optional) Disables filtering.
enable	(Optional) Enables filtering.

Command Default

The database filter is disabled.

Command Modes

Interface configuration

Area configuration

Router configuration

VRF configuration

Multi-area 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 **database-file all out** command to perform the same function that the neighbor database-filter all out, on page 1147 command performs on a neighbor basis.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to prevent flooding of OSPF LSAs to broadcast, nonbroadcast, and point-to-point networks reachable through GigabitEthernet interface 0/1/0/1:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router ospf 1
RP/0/RSP0/CPU0:router(config-ospf)# area 0
RP/0/RSP0/CPU0:router(config-ospf-ar)# interface GigabitEthernet 0/1/0/1
RP/0/RSP0/CPU0:router(config-ospf-ar-if)# database-filter all out
```

dead-interval (OSPF)

To set the interval after which a neighbor is declared dead when no hello packets are observed, use the **dead-interval** command in the appropriate mode. To return to the default time, use the **no** form of this command.

dead-interval *seconds*
no dead-interval

Syntax Description

seconds	Integer that specifies the interval (in seconds). Range is 1 to 65535. The value must be the same for all nodes on the network.

Command Default

If this command is not specified in interface configuration mode, then the interface adopts the dead interval parameter specified by the area.

If this command is not specified in area configuration mode, then the interface adopts the dead interval parameter specified for the process.

If this command is not specified at any level, then the dead interval is four times the interval set by the **hello-interval** (OSPF) command.

Command Modes

Interface configuration

Area configuration

Router configuration

Virtual-link configuration

VRF configuration

Multi-area configuration

Sham-link 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

The dead interval value must be the same for all routers and access servers on a specific network.

If the hello interval is configured, the dead interval value must be larger than the hello interval value. The dead interval value is usually configured four times larger than the hello interval value.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to set the OSPF dead interval to 40 seconds:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router ospf 1
RP/0/RSP0/CPU0:router(config-ospf)# area 0
RP/0/RSP0/CPU0:router(config-ospf-ar)# interface GigabitEthernet 0/1/0/1
RP/0/RSP0/CPU0:router(config-ospf-ar-if)# dead-interval 40
```

Related Commands

Command	Description
hello-interval (OSPF), on page 1105	Specifies the interval between hello packets that the Cisco IOS XR software sends on the interface.

default-cost (OSPF)

To specify a cost for the default summary route sent into a stub area or not-so-stubby area (NSSA), use the **default-cost** command in area configuration mode. To remove the assigned default route cost, use the **no** form of this command.

default-cost *cost*
no default-cost *cost*

Syntax Description

cost	Cost for the default summary route used for a stub or NSSA area. The acceptable value is a 24-bit number.

Command Default

cost : 1

Command Modes

Area 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **default-cost** command only on an Area Border Router (ABR) attached to a stub or an NSSA area.

In all routers and access servers attached to the stub area, the area should be configured as a stub area using the **stub** command in the area submode. Use the **default-cost** command only on an ABR attached to the stub area. The **default-cost** command provides the metric for the summary default route generated by the ABR into the stub area.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to assign a default cost of 20 to a stub area. The GigabitEthernet interface 0/4/0/3 is also configured in the stub area:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router ospf 201
RP/0/RSP0/CPU0:router(config-ospf)# area 10.15.0.0
RP/0/RSP0/CPU0:router(config-ospf-ar)# stub
RP/0/RSP0/CPU0:router(config-ospf-ar)# default-cost 20
RP/0/RSP0/CPU0:router(config-ospf-ar)# interface GigabitEthernet 0/4/0/3
```

Related Commands

Command	Description
stub (OSPF), on page 1267	Defines an area as a stub area.

default-information originate (OSPF)

To generate a default external route into an Open Shortest Path First (OSPF) routing domain, use the **default-information originate** command in the appropriate mode. To disable this feature, use the **no** form of this command.

default-information originate [**always**] [**metric** *metric-value*] [**metric-type** *type-value*] [**route-policy** *policy-name*]
no default-information originate

Syntax Description

always	(Optional) Always advertises the default route regardless of whether the routing table has a default route.
metric *metric-value*	(Optional) Specifies the metric used for generating the default route. The default metric value is 1. Range is 1 to 16777214.
metric-type *type-value*	(Optional) Specifies the external link type associated with the default route advertised into the OSPF routing domain. It can be one of the following values: **1**—Type 1 external route **2**—Type 2 external route
route-policy *policy-name*	(Optional) Specifies that a routing policy be used and the routing policy name.

Command Default

When you do not use this command in router configuration mode, no default external route is generated into an OSPF routing domain.

metric-value : 1

type-value : 2

Command Modes

Router configuration

VRF 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Whenever you use the **redistribute** or **default-information originate** command to redistribute routes into an OSPF routing domain, the software automatically becomes an Autonomous System Boundary Router (ASBR). However, an ASBR does not, by default, generate a default route into the OSPF routing domain. The software still must have a default route for itself before it generates one, except when you have specified the **always** keyword.

The **default-information originate** route-policy attach point conditionally injects the default route 0.0.0.0/0 into the OSPF link-state database, and is done by evaluating the attached policy. If any routes specified in the

policy exist in the global RIB, then the default route is inserted into the link-state database. If there is no match condition specified in the policy, the policy passes and the default route is generated into the link-state database.

For information about the default-information originate attach point, see the *OSPF Policy Attach Points* section in the *Implementing Routing Policy* chapter in *Routing Configuration Guide for Cisco ASR 9000 Series Routers*.

For information about routing policies, see the *Routing Policy Commands* chapter in the *Routing Command Reference for Cisco ASR 9000 Series Routers*.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to specify a metric of 100 for the default route redistributed into the OSPF routing domain and an external metric type of Type 1:

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#router ospf 109
RP/0/RSP0/CPU0:router(config-ospf)#redistribute igrp 108 metric 100
RP/0/RSP0/CPU0:router(config-ospf)#default-information originate metric 100 metric-type 1
```

Related Commands

Command	Description
redistribute (OSPF), on page 1180	Redistributes routes from one routing domain into a specified OSPF process.

default-metric (OSPF)

To set default metric values for routes redistributed from another protocol into the Open Shortest Path First (OSPF) protocol, use the **default-metric** command in the appropriate mode. To return to the default state, 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 16777214.

Command Default

Built-in, automatic metric translations, as appropriate for each routing protocol.

Command Modes

Router configuration

VRF 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **default-metric** command with the **redistribute** command to cause the current routing protocol to use the same metric value for all redistributed routes. A default metric helps solve the problem of redistributing routes with incompatible metrics. Whenever metrics do not convert, use a default metric to provide a reasonable substitute and enable the redistribution to proceed.

The default-metric value configured in OSPF configuration does not apply to connected routes that are redistributed to OSPF using the **redistribute connected** command. To set a non-default metric for connected routes, configure OSPF with the **redistribute connected metric** *metric-value* command.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to advertise Intermediate System-to-Intermediate System (IS-IS) protocol-derived routes into OSPF and assign a metric of 10:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router ospf 1
RP/0/RSP0/CPU0:router(config-ospf)# default-metric 10
RP/0/RSP0/CPU0:router(config-ospf)# redistribute isis IS-IS_isp
```

Related Commands

Command	Description
redistribute (OSPF), on page 1180	Redistributes routes from one routing domain into a specified OSPF process.

demand-circuit (OSPF)

To configure the Open Shortest Path First (OSPF) protocol to treat the interface as an OSPF demand circuit, use the **demand-circuit** command in the appropriate mode. To remove the demand circuit designation from the interface, use the **no** form of this command.

demand-circuit [**disable** | **enable**]
no demand-circuit

Syntax Description

disable	(Optional) Disables the interface as an OSPF demand circuit.
enable	(Optional) Enables the interface as an OSPF demand circuit.

Command Default

If this command is not specified in interface configuration mode, then the interface adopts the demand circuit parameter specified by the area.

If this command is not specified in area configuration mode, then the interface adopts the demand circuit parameter specified for the process.

If this command is not specified at any level, then the circuit is not a demand circuit.

Command Modes

Interface configuration

Area configuration

Router configuration

VRF 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

On point-to-point interfaces, only one end of the demand circuit must be configured with this command. Periodic hello messages are suppressed and periodic refreshes of link-state advertisements (LSAs) do not flood the demand circuit. Use the **demand-circuit** command to allow the underlying data link layer to be closed when the topology is stable. In point-to-multipoint topology, only the multipoint end must be configured with this command.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to set the configuration for an OSPF demand circuit:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router ospf 1
RP/0/RSP0/CPU0:router(config-ospf)# demand-circuit
```

disable-dn-bit-check

To specify that down bits should be ignored, use the **disable-dn-bit-check** command in VPN routing and forwarding (VRF) configuration mode. To specify that down bits should be considered, use the **no** form of this command.

disable-dn-bit-check
no disable-dn-bit-check

Command Default

Down bits are considered.

Command Modes

VRF configuration mode

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 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
ospf	read, write

Examples

The following example shows how to specify that down bits be ignored:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router ospf 1
RP/0/RSP0/CPU0:router(config-ospf)# vrf v1
RP/0/RSP0/CPU0:router(config-ospf-vrf)# disable-dn-bit-check
```

distance (OSPF)

To define an administrative distance, use the **distance** command in an appropriate configuration mode. To remove the **distance** command from the configuration file and restore the system to its default condition in which the software removes a distance definition, use the **no** form of this command.

distance *weight* [*ip-address wildcard-mask* [*access-list-name*]]
no distance *weight ip-address wildcard-mask* [*access-list-name*]

Syntax Description

weight	Administrative distance. Range is 10 to 255. Used alone, the *weight* argument specifies a default administrative distance that the software uses when no other specification exists for a routing information source. Routes with a distance of 255 are not installed in the routing table. lists the default administrative distances. Table 87: Default Administrative Distances, on page 1081
ip-address	(Optional) IP address in four-part, dotted-decimal notation.
wildcard-mask	(Optional) Wildcard mask in four-part, dotted decimal format. A bit set to 1 in the *mask* argument instructs the software to ignore the corresponding bit in the address value.
access-list-name	(Optional) Name of an IP access list to be applied to incoming routing updates.

Command Default

If this command is not specified, then the administrative distance is the default, as specified in Table 87: Default Administrative Distances, on page 1081.

Command Modes

Router configuration

VRF 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

An administrative distance is an integer from 10 to 255. In general, the higher the value, the lower the trust rating. An administrative distance of 255 means that the routing information source cannot be trusted at all and should be ignored. Weight values are subjective; no quantitative method exists for choosing weight values.

If an access list is used with this command, it is applied when a network is being inserted into the routing table. This behavior allows you to filter networks based on the IP prefix supplying the routing information. For example, you could filter possibly incorrect routing information from networking devices not under your administrative control.

The order in which you enter **distance** commands can affect the assigned administrative distances in unexpected ways (see the “Examples” section for further clarification).

This table lists default administrative distances.

Table 87: Default Administrative Distances

Route Source	Default Distance
Connected interface	0
Static route out on interface	0
State route to next-hop	1
EIGRP Summary Route	5
External BGP	20
Internal EIGRP	90
OSPF	110
IS-IS	115
RIP version 1 and 2	120
External EIGRP	170
Internal BGP	200
Unknown	255

Task ID

Task ID	Operations
ospf	read, write

Examples

In the following example, the **router ospf** command sets up OSPF routing instance1. The first **distance** command sets the default administrative distance to 255, which instructs the software to ignore all routing updates from networking devices for which an explicit distance has not been set. The second **distance** command sets the administrative distance for all networking devices on the Class C network 192.168.40.0 0.0.0.255 to 90.

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router ospf 1
RP/0/RSP0/CPU0:router(config-ospf)# distance 255
RP/0/RSP0/CPU0:router(config-ospf)# distance 90 192.168.40.0 0.0.0.255
```

Related Commands

Command	Description
distance bgp	Allows the use of external, internal, and local administrative distances that could be a better route to a BGP node.
distance ospf	Allows the use of external, internal, and local administrative distances that could be a better route to an OSPF node.
router ospf (OSPF), on page 1190	Configures the OSPF routing process.

distance ospf

To define Open Shortest Path First (OSPF) route administrative distances based on route type, use the **distance ospf** command in router configuration mode. To restore the default value, use the **no** form of this command.

distance ospf {**intra-area** | **inter-area** | **external**} *distance*
no distance ospf

Syntax Description

intra-area | **inter-area** | **external**	Sets the type of area. It can be one of the following values: **intra-area** —All routes within an area. **inter-area** —All routes from one area to another area. **external** —All routes from other routing domains, learned by redistribution. Any combination of the above areas is allowed.
distance	Route administrative distance.

Command Default

distance : 110

Command Modes

Router 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

You must specify one of the keywords.

Use the **distance ospf** command to perform the same function as the **distance** command used with an access list. However, the **distance ospf** command sets a distance for an entire group of routes, rather than a specific route that passes an access list.

A common reason to use the **distance ospf** command is when you have multiple OSPF processes with mutual redistribution, and you want to prefer internal routes from one over external routes from the other.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to change the external distance to 200, making the route less reliable:

```
RP/0/RSP0/CPU0:router# configure
```

```
RP/0/RSP0/CPU0:router(config)# router ospf 1
RP/0/RSP0/CPU0:router(config-ospf)# redistribute ospf 2
RP/0/RSP0/CPU0:router(config-ospf)# distance ospf external 200
RP/0/RSP0/CPU0:router(config-ospf)# exit
RP/0/RSP0/CPU0:router(config)# router ospf 2
RP/0/RSP0/CPU0:router(config-ospf)# redistribute ospf 1
RP/0/RSP0/CPU0:router(config-ospf)# distance ospf external 200
```

Related Commands

Command	Description
disable-dn-bit-check, on page 1079	Defines an administrative distance.

distribute-list

To filter networks received or transmitted in Open Shortest Path First (OSPF) updates, use the **distribute-list** command in the appropriate mode. To change or cancel the filter, use the **no** form of this command.

distribute-list {*access-list-name* {**in** | **out** [**bgp** *number* | **connected** | **ospf** *instance* | **static**]}}
no distribute-list {*access-list-name* {**in** | **out**}}

Syntax Description

access-list-name	Standard IP access list name. The list defines which networks are to be received and which are to be suppressed in routing updates.
in	Applies the access list or route-policy to incoming routing updates.
out	Applies the access list to outgoing routing updates. The **out** keyword is available only in router configuration mode.
bgp	(Optional) Applies the access list to BGP routes.
connected	(Optional) Applies the access list to connected routes.
ospf	(Optional) Applies the access list to OSPF routes (not the current OSPF process).
static	(Optional) Applies the access list to statically configured routes.

Command Default

If this command is not specified in interface configuration mode, then the interface adopts the distribute list parameter specified by the area.

If this command is not specified in area configuration mode, then the interface adopts the distribute list parameter specified for the process.

If this command is not specified at any level, then the distribute list is disabled.

Command Modes

Interface configuration

Area configuration

Router configuration

VRF configuration

Multi-area configuration

Command History

Release	Modification
Release 3.7.2	This command was introduced.
Release 4.3.1	Support was added for "if tag..." statements in **distribute-list in** *route-policy*.

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.

Use the **distribute-list** command to limit which OSPF routes are installed on this router. The **distribute-list** command does not affect the OSPF protocol itself.

The **distribute-list in** is configurable at instance (process), area, and interface levels. Regular OSPF configuration inheritance applies. Configuration is inherited from instance > area > interface levels.

The "if tag..." statements can be used in **distribute-list in** *route-policy*. The matching on route tag supports operators "eq/ge/is/le". Operator "in" is not supported.



Note When distribute-list ACL or route-policy is modified, OSPF removes all routes from the RIB, and repopulates the RIB. Due to an expected delay in the repopulation of the RIB, there could be traffic disruption.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to prevent OSPF routes from the 172.17.10.0 network from being installed if they are learned in area 0:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# ipv4 access-list 3
RP/0/RSP0/CPU0:router(config-ipv4-acl)# deny 172.17.10.0 0.0.0.255
RP/0/RSP0/CPU0:router(config-ipv4-acl)# permit any any
!
RP/0/RSP0/CPU0:router(config)# router ospf 1
RP/0/RSP0/CPU0:router(config-ospf)# area 0
RP/0/RSP0/CPU0:router(config-ospf-ar)# distribute-list 3 in
RP/0/RSP0/CPU0:router(config-ospf-ar)# interface GigabitEthernet 0/1/0/3
```

domain-id (OSPF)

To specify the Open Shortest Path First (OSPF) VPN routing and forwarding (VRF) domain ID, use the **domain-id** command in VRF configuration mode. To remove an OSPF VRF domain ID, use the **no** form of this command.

domain-id [**secondary**] **type** [**0005** | **0105** | **0205** | **8005**] **value** *value*
no domain-id [**secondary**] **type** [**0005** | **0105** | **0205** | **8005**] **value** *value*

Syntax Description

secondary	(Optional) OSPF secondary domain ID.
type	Primary OSPF domain ID in hex format.
value *value*	OSPF domain ID value in hex format (six octets).

Command Default

No domain ID is specified.

Command Modes

VRF configuration mode

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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

An OSPF domain id must be explicitly configured. The OSPF domain ID helps OSPF determine how to translate a prefix received through Border Gateway Protocol (BGP) from the remote provider edge (PE). If the domain IDs match, OSPF generates a Type 3 link state advertisement (LSA). If the domain IDs do not match, OSPF generates a Type 5 LSA.

There is only one primary domain ID. There can be multiple secondary domain IDs.



Note When an IOS XR router and an IOS router are configured as peers, the two Domain IDs must match. Manually configure the IOS XR Domain ID value to match the IOS default Domain ID value. This ensures that the routes have route code "OIA" because they are learned as inter-area routes. If the Domain IDs do not match, the routes have route code, "O-E2" because they are learned as external routes. Use the **show ip ospf** command to get the OSPF Domain ID from the IOS router. Then, set the IOS XR Domain ID to the same value using the **domain-id** command.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to specify a domain ID:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router ospf o1
RP/0/RSP0/CPU0:router(config-ospf)# vrf v1
RP/0/RSP0/CPU0:router(config-ospf-vrf)# domain-id type 0105 value AABBCCDDEEFF
```

domain-tag

To specify the Open Shortest Path First (OSPF) VPN routing and forwarding (VRF) domain tag, use the **domain-tag** command in VRF configuration mode. To remove an OSPF VRF domain tag, use the **no** form of this command.

domain-tag *tag*
no domain-tag

Syntax Description

tag	OSPF domain tag as a 32-bit value. The valid range is 0 to 4294967295.

Command Default

No OSPF VRF domain tag is specified.

Command Modes

VRF configuration mode

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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

The domain tag is added to any Type 5 link state advertisements (LSAs) generated as a result of VPN-IP routes received from Border Gateway Protocol (BGP). The domain-tag is derived from BGP autonomous system number (ASN).

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to specify the domain tag:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router ospf o1
RP/0/RSP0/CPU0:router(config-ospf)# domain-tag 234
```

exchange-timer

To set a time limit to prevent an indefinite EXCHANGE or LOADING phase during adjacency formation, use the **exchange-timer** command in router configuration mode.

exchange-timer *exchange-time-in-minutes* **hold-time** *hold-time-in-minutes* **recovery-count** *recovery-count*

Syntax Description

exchange-time-in-minutes	Duration in minutes, after which the interface is brought down if adjacencies remain in the exchange phase, in minutes. The range is from 1 to 35791394. The default value is 10.
hold-time	Specifies the option to set the time after which a recovery attempt will be made.
hold-time-in-minutes	Duration in minutes until which interface will be kept in the down state before attempting recovery, in minutes. The range is from 1 to 35791394. The default value is 5 minutes.
recovery-count	Specifies the number of attempts the router will make to bring up the adjacency towards the ASBR before keeping the adjacency shut indefinitely.
recovery-count	The number of attempts the router will make to recover the interface before the interface is permanently disabled. The recovery count value resets to the default value of 300 when the interface adjacency reaches the FULL state with the neighbor. The range is from 1 to 4294967294. The default value is 300.

Command Default

The exchange timer operates using its default values.

Command Modes

Router Configuration

Command History

Release	Modification
Release 7.10.1	This command was introduced.

Usage Guidelines

To adjust the default parameters of the exchange timer, you can use the **exchange-timer** command. We recommend configuring this command with values that suit your specific scale requirements

By default, the exchange timer is enabled with the following default values:

- Default Exchange Time: 10 minutes
- Default Hold Time: 5 minutes
- Default Recovery Count: 300

These default values have been chosen to ensure that the interface remains operational for an extended period before being permanently brought down. The total duration for attempting recovery is calculated as follows:

```
(Exchange Time + Hold Time) * Recovery Count = 3.125 days
```

Task ID

Task ID	Operation
OSPF	read, write

Example

The following example shows how to set a time limit to prevent an indefinite EXCHANGE or LOADING phase during adjacency formation.

```
Router(config)# router ospf 1
Router(config-ospf)# exchange-timer 357913 hold-time 35791 recovery-count 4294967294
```

fast-reroute (OSPFv2)

To enable IP fast reroute loop-free alternate (LFA) computation, use the **fast-reroute** command in the appropriate OSPF configuration mode. To disable the IP fast reroute loop-free alternate computation, use the **no** form of this command.

To disable loop-free alternate computation that is enabled on a higher level, use the **fast-reroute** command with **disable** keyword.

fast-reroute {**per-link** | **per-prefix**} [**disable**]
no fast-reroute

Syntax Description

per-link	Enables per-link loop-free alternate computation.
per-prefix	Enables per-prefix loop-free alternate computation.
disable	(Optional) Disables loop-free alternate computation that was enabled on a higher level.

Command Default

IP fast-reroute LFA computation is disabled.

Command Modes

Area configuration

Interface configuration

Router configuration

VRF configuration

Command History

Release	Modification
Release 4.2.0	This command was introduced and replaced the ipfrr lfa command.

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.

Only one mode of computation can be configured on an interface - per-link or per-prefix. Different modes of computations can be enabled on different interfaces; one set of interface using per-link and other set using per-prefix computation. Based on the outgoing interface of the primary path, per-link or per-prefix backup path will be computed.

Task ID

Task ID	Operation
ospf	read, write

This example shows how to enable per-link computation of loop-free alternates under interface POS 0/3/0/0:

```
RP/0/RSP0/CPU0:router(config)# router ospf 1
RP/0/RSP0/CPU0:router(config-ospf)# area 0
RP/0/RSP0/CPU0:router(config-ospf-ar)# interface GigabitEthernet 0/3/0/0
RP/0/RSP0/CPU0:router(config-ospf-ar-if)# fast-reroute per-link
```

This example shows how to enable per-prefix computation of loop-free alternates under area 0:

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#router ospf 1
RP/0/RSP0/CPU0:router(config-ospf)#area 0
RP/0/RSP0/CPU0:router(config-ospf-ar)#fast-reroute per-prefix
```

This example shows how to disable computation of loop-free alternates that was configured under area 0:

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#router ospf 1
RP/0/RSP0/CPU0:router(config-ospf)#area 0
RP/0/RSP0/CPU0:router(config-ospf-ar)#fast-reroute per-prefix
RP/0/RSP0/CPU0:router(config-ospf-ar)#interface GigabitEthernet 0/3/0/0
RP/0/RSP0/CPU0:router(config-ospf-ar-if)#fast-reroute disable
```

fast-reroute per-link exclude interface

To excludes specified interface to be used as a backup during (IPFRR) loop-free alternate (LFA) computation, use the **fast-reroute per-link exclude interface** command, in the appropriate OSPF configuration mode. To disable this feature, use the **no** form of this command.

fast-reroute per-link exclude interface *type interface-path-id*
no fast-reroute per-link exclude interface *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

No interfaces are excluded.

Command Modes

Interface configuration

Area configuration

Router configuration

VRF configuration

Command History

Release	Modification
Release 3.9.0	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
ospf	read, write

Examples

The following example shows how to exclude an interface from IP fast reroute loop-free alternate (LFA) computation:

```
RP/0/RSP0/CPU0:router(config)# router ospf 1
```

```
RP/0/RSP0/CPU0:router(config-ospf-ar-if)# fast-reroute per-link exclude interface
GigabitEthernet 0/3/2/1
```

Related Commands

Command	Description
fast-reroute (OSPFv2), on page 1092	Enables IP fast reroute loop-free alternate (LFA) computation.

fast-reroute per-prefix exclude interface (OSPFv2)

To exclude interface to be used as a backup path from fast-reroute loop-free alternate per-prefix computation, use the **fast-reroute per-prefix exclude interface** command in the appropriate OSPF configuration mode. To disable this feature, use the **no** form of this command.

fast-reroute per-prefix exclude interface *type interface-path-id*
no fast-reroute per-prefix exclude interface *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

No interfaces are excluded.

Command Modes

Interface configuration

Area configuration

Router configuration

VRF configuration

Command History

Release	Modification
Release 4.2.0	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.

Backup paths via the excluded interfaces will not be computed.

Task ID

Task ID	Operation
ospf	read, write

This example shows how to exclude interface POS0/6/0/1 from being used as a backup path:

```
RP/0/RSP0/CPU0:router#configure
```

```
RP/0/RSP0/CPU0:router(config)#router ospf 100
RP/0/RSP0/CPU0:router(config-ospf)#fast-reroute per-prefix exclude interface GigabitEthernet
 0/6/0/1
```

fast-reroute per-prefix lfa-candidate (OSPFv2)

To add interfaces to the LFA candidate list, use the **fast-reroute per-prefix lfa-candidate** command in interface configuration mode. To disable this feature, use the **no** form of this command.

fast-reroute per-prefix lfa-candidate [*interface-name*]
no fast-reroute per-prefix lfa-candidate [*interface-name*]

Syntax Description

interface-name	Specifies name of the interface to add to the LFA candidate list.

Command Default

No interfaces are added to the candidate list.

Command Modes

Interface configuration

Area configuration

Router configuration

VRF configuration

Command History

Release	Modification
Release 4.2.0	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
ospf	read, write

This example shows how to add an interface to LFA candidates:

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#router ospf 100
RP/0/RSP0/CPU0:router(config-ospf)#fast-reroute per-prefix lfa-candidate interface
GigabitEthernet 0/6/0/0
```

fast-reroute per-prefix remote-lfa (OSPFv2)

To configure fast-reroute per-prefix remote loop-free alternate (rLFA) computation for an OSPFv2 process, use the **fast-reroute per-prefix remote-lfa** command in the appropriate OSPF configuration mode. To disable this feature, use the **no** form of this command.

fast-reroute per-prefix remote-lfa {**disable** | **maximum-cost** *path-cost* | **tunnel mpls-ldp**}
no fast-reroute per-prefix remote-lfa

Syntax Description

remote-lfa	Enables remote LFA backup computation
maximum-cost *path-cost*	Sets the cost option to limit the range of remote LFAs. Range for path-cost is 1 to 4294967295.
tunnel mpls-ldp	Enables remote LFA computation using tunnel interfaces.
disable	Selectively disables remote LFA calculation under one or more areas.

Command Default

Remote LFA FRR computation is disabled.

Command Modes

Router configuration

Area configuration

Interface configuration

Command History

Release	Modification
Release 4.3.1	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.

Fast-reroute per-prefix LFA must be enabled for remote LFA to be operational. The OSPF configuration hierarchy at process level and area level are applicable for remote LFA configuration. For example, it is possible to enable remote-LFA for all OSPF areas and selectively disable (using **disable** keyword) the computation under one or more area.

Task ID

Task ID	Operation
ospf	read, write

This example shows how to enable fast-reroute per-prefix remote LFA computation for tunnel interfaces:

```
RP/0/RSP0/CPU0:router(config)#router ospf 1
```

```
RP/0/RSP0/CPU0:router(config-ospf)#fast-reroute per-prefix remote-lfa tunnel mpls-ldp
```

This example shows how to configure cost option to limit the range of remote LFAS:

```
RP/0/RSP0/CPU0:router(config)#router ospf 1
RP/0/RSP0/CPU0:router(config-ospf)#fast-reroute per-prefix remote-lfa maximum-cost 2
```

Related Commands

Command	Description
fast-reroute (OSPFv2), on page 1092	Enables IP fast reroute loop-free alternate (LFA) computation.

fast-reroute per-prefix ti-lfa

To enable per-prefix Topology Independent Loop-Free Alternate Fast Reroute (TI-LFAFRR) computation to reroute traffic around link failures, use the **fast-reroute per-prefix ti-lfa** command in the required configuration mode.

fast-reroute per-prefix ti-lfa {**enable** | **disable** }

Syntax Description

enable	Enables per-prefix Topology Independent Loop-Free Alternate Fast Reroute (TI-LFAFRR) computation to reroute traffic around link failures.
disable	Disables per-prefix Topology Independent Loop-Free Alternate Fast Reroute (TI-LFAFRR) computation.

Command Default

TI-LFAFRR computation is not enabled by default.

Command Modes

Interface configuration

Area configuration

Router configuration

Command History

Release	Modification
Release 5.3.0	This command was introduced.

Usage Guidelines

OSPFv2 Loop-Free Alternate Fast Reroute TI-LFA is built on top of per-prefix LFA and as such requires per-prefix LFA to be enabled on the interface, where TI-LFA is enabled.

Task ID

Task ID	Operation
ospf	read, write

This example shows how to enable per-prefix Topology Independent Loop-Free Alternate Fast Reroute computation for the interface:

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#router ospf
RP/0/RSP0/CPU0:router(config-ospf)#area 1
RP/0/RSP0/CPU0:router(config-ospf-ar)#interface
GigabitEternet0/5/0/0 fast-reroute per-prefix ti-lfa enable
```

fast-reroute per-prefix use-candidate-only (OSPFv2)

To restrict the backup interfaces to those that are present on the LFA candidate list, use the **fast-reroute per-prefix use-candidate-only** command in router OSPF configuration mode. To disable this feature, use the **no** form of this command.

fast-reroute per-prefix use-candidate-only [**enable** | **disable**]
fast-reroute per-prefix use-candidate-only

Syntax Description

enable	Enables backup selection from candidate-list only.
disable	Disables backup selection from candidate-list only.

Command Default

Disabled.

Command Modes

Router OSPF configuration

Command History

Release	Modification
Release 4.2.0	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
ospf	read, write

This example shows how to restrict the backup interfaces to those that are present on the LFA candidate list:

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#router ospf 100
RP/0/RSP0/CPU0:router(config-ospf)#fast-reroute per-prefix use-candidate-only
```

flood-reduction (OSPF)

To suppress the unnecessary flooding of link-state advertisements (LSAs) in stable topologies, use the **flood-reduction** command in the appropriate mode. To remove this functionality from the configuration, use the **no** form of this command.

flood-reduction [**enable** | **disable**]
no flood-reduction [**enable** | **disable**]

Syntax Description

enable	(Optional) Turns on this functionality at a specific level.
disable	(Optional) Turns off this functionality at a specific level.

Command Default

If this command is not specified in interface configuration mode, then the interface adopts the flood reduction parameter specified by the area.

If this command is not specified in area configuration mode, then the interface adopts the flood reduction parameter specified for the process.

If this command is not specified at any level, then flood reduction is disabled.

Command Modes

Interface configuration

Area configuration

Router configuration

VRF 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

All routers supporting the OSPF demand circuit are compatible and can interact with routers supporting flooding reduction.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to reduce the flooding of unnecessary LSAs for area 0:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router ospf 1
```

```
RP/0/RSP0/CPU0:router(config-ospf)# area 0
RP/0/RSP0/CPU0:router(config-ospf-ar)# interface GigabitEthernet 0/1/0/3
RP/0/RSP0/CPU0:router(config-ospf-ar-if)# flood-reduction
```

Related Commands

Command	Description
show ospf interface, on page 1221	Displays OSPF-related interface information.
show ospf neighbor, on page 1232	Displays OSPF neighbor information on an individual interface basis.

hello-interval (OSPF)

To specify the interval between consecutive hello packets that are sent on the Open Shortest Path First (OSPF) interface, use the **hello-interval** command in the appropriate mode. To return to the default time, use the **no** form of this command.

hello-interval *seconds*
no hello-interval

Syntax Description

seconds	Interval (in seconds). The value must be the same for all nodes on a specific network. Range is 1 to 65535.

Command Default

If this command is not specified in interface configuration mode, then the interface adopts the hello interval parameter specified by the area.

If this command is not specified in area configuration mode, then the interface adopts the hello interval parameter specified for the process.

If this command is not specified at any level, then the hello interval is 10 seconds (broadcast) or 30 seconds (non-broadcast).

Command Modes

Interface configuration

Area configuration

Router configuration

Virtual-link configuration

VRF configuration

Multi-area configuration

Sham-link 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

The hello interval value is advertised in the hello packets. The shorter the hello interval, the faster topological changes are detected, but more routing traffic occurs. This value must be the same for all routers and access servers on a specific network.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to set the interval between hello packets to 15 seconds:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router ospf 1
RP/0/RSP0/CPU0:router(config-ospf)# area 0
RP/0/RSP0/CPU0:router(config-ospf-ar)# interface GigabitEthernet 0/1/0/1
RP/0/RSP0/CPU0:router(config-ospf-ar-if)# hello-interval 15
```

Related Commands

Command	Description
dead-interval (OSPF), on page 1069	Sets the time period for which hello packets are suspended before neighbors declare the router down.

ignore lsa mospf

To suppress the sending of syslog messages when the router receives link-state advertisement (LSA) Type 6 multicast Open Shortest Path First (MOSPF) packets, which are unsupported, use the **ignore lsa mospf** command in an appropriate configuration mode. To restore the sending of syslog messages, use the **no** form of this command.

ignore lsa mospf
no ignore lsa mospf

This command has no keywords or arguments.

Command Default

When you do not specify this command in router configuration mode, each MOSPF packet received by the router causes the router to send a syslog message.

Command Modes

Router configuration

VRF 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Cisco routers do not support LSA Type 6 (MOSPF), and they generate syslog messages if they receive such packets. If the router is receiving many MOSPF packets, you might want to configure the router to ignore the packets and thus prevent a large number of syslog messages.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to configure the router to suppress the sending of syslog messages when it receives MOSPF packets:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router ospf 1
RP/0/RSP0/CPU0:router(config-ospf)# ignore lsa mospf
```

interface (OSPF)

To define the interfaces on which the Open Shortest Path First (OSPF) protocol runs, use the **interface** command in area configuration mode. To disable OSPF routing for interfaces, use the **interface** form of this command.

interface *type interface-path-id*
no interface *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 virtualinterface. **Note** se 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

When you do not specify this command in configuration mode, OSPF routing for interfaces is not enabled.

Command Modes

Area 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **interface** command to associate a specific interface with an area. The interface remains associated with the area even when the IP address of the interface changes.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how the OSPF routing process 109 defines four OSPF areas (0, 2, 3, and 10.9.50.0), and associates an interface with each area:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router ospf 109
RP/0/RSP0/CPU0:router(config-ospf)# area 0
RP/0/RSP0/CPU0:router(config-ospf-ar)# interface GigabitEthernet 4/0/0/3
!
```

```
RP/0/RSP0/CPU0:router(config-ospf)# area 2
RP/0/RSP0/CPU0:router(config-ospf-ar)# interface GigabitEthernet 0/1/0/3
!
RP/0/RSP0/CPU0:router(config-ospf)# area 3
RP/0/RSP0/CPU0:router(config-ospf-ar)# interface GigabitEthernet 3/0/0/2
!
RP/0/RSP0/CPU0:router(config-ospf)# area 10.9.50.0
RP/0/RSP0/CPU0:router(config-ospf-ar)# interface GigabitEthernet 3/0/0/1
```

log adjacency changes (OSPF)

To configure the router to send a syslog message when the state of an Open Shortest Path First (OSPF) neighbor changes, use the **log adjacency changes** command in router configuration mode. To turn off this function, use the **disable** keyword. To log all state changes, use the **detail** keyword.

log adjacency changes {**detail** | **disable**}

Syntax Description

detail	Provides all (DOWN, INIT, 2WAY, EXSTART, EXCHANGE, LOADING, FULL) adjacency state changes.
disable	Disables sending adjacency change messages.

Command Default

The router sends a syslog message when the state of an OSPF neighbor changes.

Command Modes

Router 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **log adjacency changes** command to display high-level changes to the state of the peer relationship. Configure this command if you want to know about OSPF neighbor changes.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to configure the software to send a syslog message for any OSPF neighbor state changes:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router ospf 109
RP/0/RSP0/CPU0:router(config-ospf)# log adjacency changes detail
```

loopback stub-network

To enable advertising loopback as stub networks, use the **loopback stub-network** command in an appropriate configuration mode. To disable advertising loopback as stubnetworks, use the **no** form of this command.

loopback stub-network [**enable** | **disable**]
no loopback stub-network

Syntax Description

enable	(Optional) Enables advertising loopbacks as stub networks.
disable	(Optional) Disables advertising loopbacks as stub networks.

Command Default

By default, OSPF advertises loopbacks as stub hosts.

Command Modes

OSPF interface configuration

OSPF router configuration

OSPF area configuration

Command History

Release	Modification
Release 3.9.0	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.

In the interface submode, the command can be enabled only on loopback interfaces.

Task ID

Task ID	Operation
ospf	read, write

Examples

The following example shows how to enable advertising loopback as a stub network, under OSPF interface configuration:

```
RP/0/RSP0/CPU0:router(config)#router ospf 100
RP/0/RSP0/CPU0:router(config-ospf)#loopback stub-network enable
```

Related Commands

Command	Description
show ospf interface, on page 1221	Displays Open Shortest Path First (OSPF) interface information.

lpts punt excessive-flow-trap penalty-timeout ospf

To set the penalty timeout for the OSPF protocol, use the **lpts punt excessive-flow-trap penalty-timeout ospf** command in the Global Configuration mode. To restore the default penalty timeout value, use the **no** form of this command.

lpts punt excessive-flow-trap {**penalty-timeout ospf**}*timeout*
no lpts punt excessive-flow-trap {**penalty-timeout ospf**}

Syntax Description

timeout	The penalty timeout value for the ospf protocol in minutes. It is the period of time at which ospf bad flow remains to be in bad actor state. Value ranges from 1 to 1000.

Command Default

The default penalty timeout value is 15 minutes.

Command Modes

Global Configuration mode

Command History

Release	Modification
Release 6.0.1	This command was introduced.

Usage Guidelines

You must be in a user group associated with a task group that includes the proper task IDs. The command reference guides include the task IDs required for each command. If you suspect user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

If the penalty-timeout value for ospf is configured as 20, then ospf packets are dropped for 20 minutes.

Task ID

Task ID	Operations
lpts	read
basic-services	read-write

Examples

This example shows to set penalty time out for OSPF bad actor:

```
RP/0/RSP0/CPU0:router(config)# lpts punt excessive-flow-trap
RP/0/RSP0/CPU0:router(config-control-plane-policer)#  penalty-timeout ospf <1-1000>
```

Related Commands

Command	Description
show running-config lpts punt excessive-flow-trap, on page 537	Displays the running configuration for the Excessive Punt Flow Trap feature.

max external-lsa

To specify the maximum number of external LSAs that the router will accept from a neighbor, use the **max-external-lsa** command in router configuration mode.

max-external-lsa *lsa-limit* *logging-threshold-percentage* [**suppress-neighbor** | **warning-only**]

Syntax Description

lsa-limit	Specifies the maximum number of external LSAs that a router will accept from a neighbor. If the neighbor sends more than the specified number of LSAs, the adjacency will be brought down and kept down for a specified interval. The range is from 1 to 4294967294. The default value is 50,000.
logging-threshold-percentage	Specifies the percentage of the total number of external LSAs from an ASBR at which a warning log is given. The default value is 75.
suppress-neighbor	Specifies that the link towards an ASBR be brought down if the ASBR is the immediate neighbor. This option is enabled by default.
warning-only	Specifies the option to overwrite the suppress-neighbor behavior. If the **warning-only** keyword is configured instead of **suppress-neighbor** keyword, the link to the ASBR is not brought down when the limit of external LSAs is reached. Only alert logs is given.

Command Default

If the command is not configured, the ASBR does not monitor unsuccessful adjacencies towards other ASBRs caused by the ASBR reaching the configured maximum limit for external Link State Advertisements (LSAs). Additionally, the ASBR does not incorporate a counter that, upon reaching the specified number of attempts to establish adjacency, causes the neighbor to cease further efforts, triggering permanent adjacency downtime.

Command Modes

Router Configuration

Command History

Release	Modification
Release 7.10.1	This command was introduced.

Usage Guidelines

We recommend configuring this command with values that suit your specific scale requirements.

Task ID

Task ID	Operation
OSPF	read, write

Example

The following example show how to configure the maximum number of external LSAs, the warning log percentage, and the link-down option for an OSPF router.

```
Router(config)# router ospf 1
Router(config-ospf)# area 0
Router(config-ospf-ar)# max-external-lsa 50000 suppress-neighbor
```

max-lsa

To limit the number of nonself-generated link-state advertisements (LSAs) that an Open Shortest Path First (OSPF) routing process can keep in the OSPF link-state database (LSDB), use the **max-lsa** command in router configuration mode. To remove the limit of non self-generated LSAs that an OSPF routing process can keep in the OSPF LSDB, use the **no** form of this command.

max-lsa *max* [*threshold*] [**warning-only**] [**ignore-time** *value*] [**ignore-count** *value*] [**reset-time** *value*]
no max-lsa *max* [*threshold*] [**warning-only**] [**ignore-time** *value*] [**ignore-count** *value*] [**reset-time** *value*]

Syntax Description

max	Maximum number of nonself-generated LSAs the OSPF process can keep in the OSPF LSBD.
threshold	(Optional) The percentage of the maximum LSA number, as specified by the maximum-number argument, at which a warning message is logged. The default is 75 percent.
warning-only	(Optional) Specifies that only a warning message is sent when the maximum limit for LSAs is exceeded. Disabled by default.
ignore-time *value*	(Optional) Specifies the time, in minutes, to ignore all neighbors after the maximum limit of LSAs has been exceeded. The default is 5 minutes.
ignore-count *value*	(Optional) Specifies the number of times the OSPF process can consecutively be placed into the ignore state. The default is 5 times.
reset-time *value*	(Optional) Specifies the time, in minutes, after which the ignore count is reset to zero. The default is 2 times **ignore-time** .

Command Default

Disabled

Command Modes

Router configuration

VRF 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

This command allows you to protect the OSPF routing process from the large number of received LSAs that can result from a misconfiguration on another router in the OSPF domain (for example, the redistribution of a large number of IP prefixes to OSPF).

When this feature is enabled, the router keeps count of the number of all received (nonself-generated) LSAs. When the configured *threshold* value is reached, an error message is logged. When the configured *max* number of received LSAs is exceeded, the router stops accepting new LSAs.

If the count of received LSAs is higher than the configured *max* number after one minute, the OSPF process disables all adjacencies in the given context and clears the OSPF database. This state is called the ignore state. In this state, all OSPF packets received on all interfaces belonging to the OSPF instance are ignored and no OSPF packets are generated on its interfaces. The OSPF process remains in the ignore state for the duration of the configured **ignore-time** . When the **ignore-time** expires, the OSPF process returns to normal operation and starts building adjacencies on all its interfaces.

To prevent the OSPF instance from endlessly oscillating between its normal state and the ignore state, as a result of the LSA count immediately exceeding the *max* number again after it returns from the ignore state, the OSPF instance keeps a count of how many times it has been in the ignore state. This counter is called the **ignore-count** . If the **ignore-count** exceeds its configured value, the OSPF instance remains in the ignore state permanently.

To return the OSPF instance to its normal state, you must issue the **clear ip ospf** command. The **ignore-count** is reset to zero if the LSA count does not exceed the *max* number again during the time configured by the **reset-time** keyword.

If you use the **warning-only** keyword, the OSPF instance never enters the ignore state. When LSA count exceeds the *max* number, the OSPF process logs an error message and the OSPF instance continues in its normal state operation.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to configure the OSPF instance to accept 12000 nonself-generated LSAs in the global routing table, and 1000 nonself-generated LSAs in VRF V1.

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router ospf 0
RP/0/RSP0/CPU0:router(config-ospf)# max-lsa 12000
RP/0/RSP0/CPU0:router(config-ospf)# vrf V1
RP/0/RSP0/CPU0:router(config-ospf)# max-lsa 1000
```

The following example shows how to display the current status of the OSPF instance:

```
RP/0/RSP0/CPU0:router# show ospf 0

  Routing Process "ospf 0" with ID 10.0.0.2
  NSR (Non-stop routing) is Disabled
  Supports only single TOS(TOS0) routes
  Supports opaque LSA
  It is an area border router
  Maximum number of non self-generated LSA allowed 12000
     Current number of non self-generated LSA 1
     Threshold for warning message 75%
     Ignore-time 5 minutes, reset-time 10 minutes
     Ignore-count allowed 5, current ignore-count 0
```

Related Commands

Command	Description
show ospf, on page 1199	Displays general information about Open Shortest Path First (OSPF) routing processes.

max-metric

To configure the Open Shortest Path First (OSPF) protocol to signal other networking devices not to prefer the local router as an intermediate hop in their shortest path first (SPF) calculations, use the **max-metric** command in router configuration mode. To disable this function, use the **no** form of this command.

max-metric router-lsa [**external-lsa** *overriding metric*] [**include-stub**] [**on-proc-migration**] [**on-proc-restart**] [**on-startup**] [**on-switchover**] [**wait-for-bgp**] [**summary-lsa**]
no max-metric router-lsa

Syntax Description

router-lsa	Always originates router link-state advertisements (LSAs) with the maximum metric.
external-lsa *overriding metric*	(Optional) Overrides the external-lsa metric with the max-metric value.The *overriding metric* argument specifies the number of in-summary-LSAs. The range is 1 to 16777215>. The default is 16711680.
include-stub	(Optional) Advertises stub links in router-LSA with the max-metric value (0xFFFF).
on-proc-migration *time*	(Optional) Sets the maximum metric temporarily after a process migration to originate router-LSAs with the max-metric value. The *time* range is 5 to 86400 seconds.
on-proc-restart *time*	(Optional) Sets the maximum metric temporarily after a process restart to originate router-LSAs with the max-metric value. The *time* range is 5 to 86400 seconds.
on-startup *time*	(Optional) Sets the maximum metric temporarily after a reboot to originate router-LSAs with the max-metric value. The *time* range is 5 to 86400 seconds.
on-switchover *time*	(Optional) Sets the maximum metric temporarily after a switchover to originate router-LSAs with the max-metric value. The *time* range is 5 to 86400 seconds. **Note** OSPF will not populate maximum metric on the router's generated LSAs, when the OSPF routing process is configured to support Nonstop Routing (NSR) or Nonstop Forwarding/Graceful restart (NSF/GR).
wait-for-bgp	(Optional) Causes OSPF to originate router LSAs with the maximum metric and allows Border Gateway Protocol (BGP) to decide when to start originating router LSAs with a normal metric instead of the maximum metric.
summary-lsa	(Optional) specifies the number of in summary-LSAs. The range is 1 to 16777215. The default is 16711680.

Command Default

Router LSAs are originated with normal link metrics.

overriding-metric :16711680

Command Modes

Router configuration

VRF 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **max-metric** command to cause the software to originate router LSAs with router link metrics set to LSInfinity (0XFFFF). This feature can be useful in Internet backbone routers that run both OSPF and BGP because OSPF converges more quickly than BGP and may begin attracting traffic before BGP has converged, resulting in dropped traffic.

If this command is configured, the router advertises its locally generated router LSAs with a metric of 0XFFFF. This action allows the router to converge but not attract transit traffic if there are better, alternative paths around this router. After the specified *announce-time* value or notification from BGP has expired, the router advertises the local router LSAs with the normal metric (interface cost).

If this command is configured with the **on-startup** keyword, then the maximum metric is temporarily set only after reboot is initiated. If this command is configured without the **on-startup** keyword, then the maximum metric is permanently used until the configuration is removed.

If the **include-stub** keyword is enabled, the stub-links in the router LSA will be sent with the max-metric. If the **summary-lsa** keyword is enabled, all self-generated summary LSAs will have a metric set to 0xFF0000, unless the metric value is specified with the max-metric value parameter. If the **external-lsa** keyword is enabled, all self-generated external LSAs will have a metric set to 0xFF0000, unless the metric value is specified with the max-metric value parameter.

This command might be useful when you want to connect a router to an OSPF network, but do not want real traffic flowing through it if there are better, alternative paths. If there are no alternative paths, this router still accepts transit traffic as before.

Some cases where this command might be useful are as follows:

- During a router reload, you prefer that OSPF wait for BGP to converge before accepting transit traffic. If there are no alternative paths, the router still accepts transit traffic.
- A router is in critical condition (for example, it has a very high CPU load or does not have enough memory to store all LSAs or build the routing table).
- When you want to gracefully introduce or remove a router to or from the network.
- When you have a test router in a lab, connected to a production network.



Note For older OSPF implementations (RFC 1247), router links in received router LSAs with a metric and cost of LSInfinity are not used during SPF calculations. Hence, no transit traffic is set to the routers originating such router LSAs.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to configure OSPF to originate router LSAs with the maximum metric until BGP indicates that it has converged:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router ospf 109
RP/0/RSP0/CPU0:router(config-ospf)# max-metric router-lsa on-startup wait-for-bgp
```

maximum interfaces (OSPF)

To limit the number of interfaces that can be configured for an Open Shortest Path First (OSPF) process, use the **maximum interfaces** command in the appropriate mode. To return to the default limit, use the **no** form of this command.

maximum interfaces *number-interfaces*
no maximum interfaces

Syntax Description

number-interfaces	Number of interfaces. Range is 1 to 4294967295.

Command Default

If the command is not specified, the default is 1024.

Command Modes

Router configuration

VRF configuration

Command History

Release	Modification
Release 3.7.2	This command was introduced.
Release 4.1.1	The range for number of interfaces was changed to 1 to 4294967295 from 1 to 1024. The default number of interfaces was changed to 1024 from 255.

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.

Use the **maximum interface** command to increase or decrease the limit on the number of interfaces configured for an OSPF process.

You cannot configure a limit lower than the number of interfaces currently configured for the OSPF process. To lower the limit, remove interfaces from the OSPF configuration until the number of configured interfaces is at or below the desired limit. You may then apply the new, lower limit.

Task ID

Task ID	Operations
ospf	read, write

Examples

This example shows how to configure a maximum interface limit of 1500 on a router:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router ospf 109
RP/0/RSP0/CPU0:router(config-ospf)# maximum interfaces 1500
```

Related Commands

Command	Description
show ospf interface, on page 1221	Displays OSPF interface information.

maximum paths (OSPF)

To control the maximum number of parallel routes that the Open Shortest Path First (OSPF) protocol can support, use the **maximum paths** command in an appropriate configuration mode. To remove the **maximum paths** command from the configuration file and restore the system to its default condition with respect to the routing protocol, use the **no** form of this command.

maximum paths *maximum-routes-number*
no maximum paths

Syntax Description

maximum-routes-number	Maximum number of parallel routes that OSPF can install in a routing table. Range is 1 to 64. **Note** The maximum number of paths that can be configured is 64.

Command Default

The default value for maximum-paths depends on the platform supported maximum-paths value.

Command Modes

Router configuration

VRF configuration

Command History

Release	Modification
Release 3.7.2	This command was introduced.
Release 5.3.0	ECMP support extended from 32 to 64 paths.

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.

The default value for maximum-path depends on the platform supported maximum-path value. Any custom value you define for the maximum-path parameter must be within the maximum value supported by platform. The configuration will be rejected if the value you have specified is more then what the platform supports.

When the maximum number of parallel routes is reduced, all existing paths are pruned and paths reinstalled at the new maximum number. During this route-reduction period, you may experience some packet loss for a few seconds. This may impact route traffic.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to allow a maximum of two paths to a destination:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router ospf 109
RP/0/RSP0/CPU0:router(config-ospf)# maximum paths 2
```

maximum redistributed-prefixes (OSPF)

To limit the aggregate number of prefixes that can be redistributed into an Open Shortest Path First (OSPF) process, use the **maximum redistributed-prefix** command in the appropriate mode. To return to the default limit, use the **no** form of this command.

maximum redistributed-prefixes *maximum* [*threshold-value*] [**warning-only**]
no maximum redistributed-prefixes

Syntax Description

maximum	Number of routes. Range is 1 to 4294967295.
threshold-value	(Optional) Threshold value (as a percentage) at which to generate a warning message. Range is 1 to 100.
warning-only	(Optional) Gives only a warning when the limit is exceeded.

Command Default

If the command is not specified, the default is 10000.

The threshold value defaults to 75 percent.

Command Modes

Router configuration

VRF configuration

Command History

Release	Modification
Release 3.7.2	This command was introduced.
Release 7.9.1	• Support for the maximum redistributed-prefixes limit was added for the prefixes redistributed as Type-3 LSAs. • Support for counting only the number of redistributed prefixes instead of the number of LSAs was added. (If multiple prefixes are aggregated in an LSA using the summary-prefix command, individual prefixes will be counted)

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.
- Use the **maximum redistributed-prefixes** command to increase or decrease the maximum number of prefixes (also referred to as routes) redistributed for an OSPF process.

  If the maximum value specified is less than the existing number of redistributed prefixes, then the prefixes are reset and the redistribution is restarted.
- Starting Cisco IOS XR Software Release 7.9.1, when the **summary-prefix** command is used to create aggregate LSAs for the routes redistributed into OSPF, each individual redistributed prefix mapped under the subnet advertised in the LSA is counted instead of the number of LSAs. For example, if 5 redistributed prefixes are aggregated by one LSA using the **summary-prefix** command, only the number of redistributed prefixes, which is 5 in this case, is counted and not the number of LSAs, which is 1.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to configure a maximum number of routes that can be redistributed for an OSPF routing process:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router ospf 109
RP/0/RSP0/CPU0:router(config-ospf)# maximum redistributed-prefixes 15000
```

Related Commands

Command	Description
show ospf routes, on page 1246	Displays the OSPF topology table.

message-digest-key

To specify a key used with Open Shortest Path First (OSPF) Message Digest 5 (MD5) authentication, use the **message-digest-key** command in the appropriate mode. To remove an old MD5 key, use the **no** form of this command.

message-digest-key *key-id* **md5** {*key* | **clear** *key* | **encrypted** *key*}
no message-digest-key *key-id*

Syntax Description

key-id	Key number. Range is 1 to 255.
md5	Enables OSPF MD5 authentication.
key	Alphanumeric string of up to 16 characters.
clear	Specifies that the key be clear text.
encrypted	Specifies that the key be encrypted using a two-way algorithm.

Command Default

If this command is not specified in interface configuration mode, then the interface adopts the message digest key parameter specified by the area.

If this command is not specified in area configuration mode, then the interface adopts the message digest key parameter specified for the process.

If this command is not specified at any level, then OSPF MD5 authentication is disabled.

Command Modes

Interface configuration

Area configuration

Router configuration

Virtual-link configuration

VRF configuration

Multi-area configuration

Sham-link 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Usually, one key individual interface is used to generate authentication information when packets are sent and to authenticate incoming packets. The same key identifier on the neighbor router must have the same *key* value.

For authentication to be enabled, you must configure the **message-digest-key** command together with the **authentication** command and its **message-digest** keyword. Both the **message-digest-key** and **authentication** commands can be inherited from a higher configuration level.

The process of changing keys is as follows. Suppose the current configuration is:

```
interface GigabitEthernet 0/3/0/2
 message-digest-key 100 md5 OLD
```

You change the configuration to the following:

```
interface GigabitEthernet 0/3/0/2
 message-digest-key 101 md5 NEW
```

The system assumes its neighbors do not have the new key yet, so it begins a rollover process. It sends multiple copies of the same packet, each authenticated by different keys. In this example, the system sends out two copies of the same packet—the first one authenticated by key 100 and the second one authenticated by key 101.

Rollover allows neighboring routers to continue communication while the network administrator is updating them with the new key. Rollover stops after the local system finds that all its neighbors know the new key. The system detects that a neighbor has the new key when it receives packets from the neighbor authenticated by the new key.

After all neighbors have been updated with the new key, the old key should be removed. In this example, you would enter the following:

```
interface ethernet 1
 no ospf message-digest-key 100
```

Then, only key 101 is used for authentication on interface 1.

We recommend that you not keep more than one key individual interface. Every time you add a new key, you should remove the old key to prevent the local system from continuing to communicate with a hostile system that knows the old key. Removing the old key also reduces overhead during rollover.

Note The MD5 key is always stored in encrypted format on the router. The **clear** and **encrypted** keywords inform the router whether the value that is entered is encrypted or unencrypted.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to set a new key 19 with the password *8ry4222* :

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router ospf 109
RP/0/RSP0/CPU0:router(config-ospf)# area 0

RP/0/RSP0/CPU0:router(config-ospf-ar)# interface GigabitEthernet 0/1/0/1
RP/0/RSP0/CPU0:router(config-ospf-ar-if)# message-digest-key 19 md5 8ry4222
```

Related Commands

Command	Description
area (OSPF), on page 1047	Configures an OSPF area.
authentication (OSPF), on page 1049	Enables plain text, MD5 authentication, or null authentication for an OSPF interface.
default-cost (OSPF), on page 1071	Enables authentication for an OSPF area.

microloop avoidance (OSPF)

To avoid micro-loops by delaying the convergence of all or protected prefixes, use the **microloop avoidance** command.

To disable this function, use the **no** prefix for this command.

microloop avoidance [**protected** | **rib-update-delay** *delay*]
no microloop avoidance

Syntax Description

(none)	Delays convergence of all prefixes.
protected	(Optional) Delays convergence of protected prefixes..
rib-update-delay *delay*	(Optional) Delays convergence of all prefixes and updates RIB after the configured delay. The range is 1 to 60000 milliseconds.

Command Default

Microloop avoidance is disabled by default.

Command Modes

router ospf configuration

Command History

Release	Modification
Release 6.1.1	This command was introduced.

Usage Guidelines

When the network converges after a link failure restoration, micro-loops can form due to inconsistencies in the forwarding tables of different routers. By delaying the convergence of prefixes, you can avoid the formation of micro-loops.

You can delay the convergence of all or protected prefixes by using the **microloop avoidance** command. When configured, the command applies to all prefixes by default. To enable it for only protected prefixes, use the **protected** option.

You can delay updates to the RIB, by using the **rib-update-delay** option.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to configure microloop avoidance with OSPF:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router ospf 100
RP/0/RSP0/CPU0:router(config-ospf)# microloop avoidance protected
```

mpls ldp auto-config (OSPF)

To enable Label Distribution Protocol (LDP)-Interior Gateway Protocol (IGP) interface automatic configuration, use the **mpls ldp auto-config** command in the appropriate mode. To disable LDP-IGP interface automatic configuration, use the **no** form of this command.

mpls ldp auto-config
no mpls ldp auto-config

Command Default

LDP-IGP interface automatic configuration is disabled for OSPF.

Command Modes

Interface configuration

Area configuration

Router 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 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
ospf	read, write

Examples

The following example shows how to enable LDP-IGP interface automatic configuration:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router ospf o1
RP/0/RSP0/CPU0:router(config-ospf)# mpls ldp auto-config
```

mpls ldp sync (OSPF)

To enable Label Distribution Protocol (LDP)-Interior Gateway Protocol (IGP) synchronization, use the **mpls ldp sync** command in the appropriate mode. To disable LDP-IGP synchronization, use the **no** form of this command.

mpls ldp sync [**disable**]
no mpls ldp sync

Syntax Description

disable	(Optional) Disables MPLS LDP synchronization from within the OSPF interface and area configuration submodes only. For the OSPF router configuration mode, use the **no** form of the command.

Command Default

LDP-IGP synchronization is disabled for OSPF.

Command Modes

Interface configuration

Area configuration

Router 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 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
ospf	read, write

Examples

The following example shows how to enable LDP-IGP synchronization:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router ospf o1
RP/0/RSP0/CPU0:router(config-ospf)# mpls ldp sync
```

mpls traffic-eng (OSPF)

To configure an Open Shortest Path First (OSPF) area for Multiprotocol Label Switching traffic engineering (MPLS TE), use the **mpls traffic-eng** command in the appropriate configuration mode. To remove the MPLS TE from an area, use the **no** form of this command.

mpls traffic-eng
no mpls traffic-eng

Syntax Description

This command has no keywords or arguments.

Command Default

MPLS TE is not configured for OSPF.

Command Modes

Area configuration

VRF 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

You must configure the **mpls traffic-eng** command for OSPF to support MPLS traffic engineering. OSPF provides the flooding mechanism that is used to flood TE link information.

Note This command is supported only in the default VRF mode.

We recommend that you configure the **mpls traffic-eng router-id** command instead of using the **router-id** command in global configuration mode.

OSPF support for MPLS TE is a component of the overall MPLS TE feature. Other MPLS TE software components must also be configured for this feature to be fully supported.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to associate loopback interface 0 with area 0, and area 0 is declared to be an MPLS area:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router ospf 1
```

```
RP/0/RSP0/CPU0:router(config-ospf)# router-id 10.10.10.10
RP/0/RSP0/CPU0:router(config-ospf)# mpls traffic-eng router-id loopback 0
RP/0/RSP0/CPU0:router(config-ospf)# area 0
RP/0/RSP0/CPU0:router(config-ospf)# mpls traffic-eng
RP/0/RSP0/CPU0:router(config-ospf-ar)# interface loopback 0
```

Related Commands

Command	Description
capability opaque disable, on page 1055	Controls the OSPF opaque LSA support capability.
mpls traffic-eng multicast-intact (OSPF), on page 1137	Specifies that the traffic engineering router identifier for the node is the IP address associated with a given interface.
router-id (OSPF), on page 1188	Configures a router ID for the OSPF process.

mpls traffic-eng igp-intact (OSPF)

To ensure that the OSPF protocol installs at least one IPv4 next-hop when it adds the tunnel next-hops (igp-shortcuts), use the **mpls traffic-eng igp-intact** command in the router configuration mode. To disable IGP-intact, use the **no** form of this command.

mpls traffic-eng igp-intact
no mpls traffic-eng igp-intact

Command Default

IGP-intact is disabled.

Command Modes

Router 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

The OSPF protocol adds both tunnel next-hops (igp-shortcuts) and IPv4 next-hops to the next-hop list in the Routing Information Base (RIB), until it reaches the maximum number of paths. When IGP-intact is enabled, it ensures that the Routing Information Base (RIB) always has at least one IPv4 next-hop present in the list of next-hops when the number of paths is at maximum.

Note IGP-intact should be used only when Policy-Based Tunnel Selection (PBTS) is in use.

Configure IGP-intact only when Policy-Based Tunnel Selection (PBTS) is in use. This ensures that at least one IPv4 next-hop is available for the default Differentiated Services Code Point (DSCP) traffic class. It also prevents traffic loss for other DSCP traffic classes by diverting such traffic to an IPv4 next-hop when the corresponding tunnel is unavailable for forwarding.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to enable IGP-intact:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router ospf 1
RP/0/RSP0/CPU0:router(config-ospf)# mpls traffic-eng igp-intact
```

Related Commands

Command	Description
maximum paths (OSPF), on page 1123	Configures the maximum number of parallel routes that the Open Shortest Path First (OSPF) protocol can support.

mpls traffic-eng multicast-intact (OSPF)

To enable multicast-intact for Open Shortest Path First (OSPF) routes so that multicast-intact paths are published to the Routing Information Base (RIB), use the **mpls traffic-eng multicast-intact** command in the appropriate configuration mode. To remove the MPLS TE area, use the **no** form of this command.

mpls traffic-eng multicast-intact
no mpls traffic-eng multicast-intact

Syntax Description

This command has no keywords or arguments.

Command Default

MPLS TE is not configured for OSPF.

Command Modes

Router configuration

VRF 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

OSPF support for MPLS TE is a component of the overall MPLS TE feature. Other MPLS TE software components must also be configured for this feature to be fully supported.

Note This command is supported only in the default VRF mode.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to enable publishing of multicast-intact paths to RIB:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router ospf 1
RP/0/RSP0/CPU0:router(config-ospf)# mpls traffic-eng multicast-intact
```

mpls traffic-eng ldp-sync-update (OSPF)

To enable Label Distribution Protocol (LDP)-OSPF Type 1 and Type 10 link-state advertisement (LSA) synchronization, use the **mpls traffic-eng ldp-sync-update** command in the appropriate mode. To disable LDP-LSA synchronization, use the **no** form of this command.

mpls traffic-eng ldp-sync-update
no mpls traffic-eng ldp-sync-update

Command Default

LDP-LSA synchronization is disabled for OSPF.

Command Modes

Router configuration

Interface configuration

Area configuration

Command History

Release	Modification
Release 4.2.4	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.

Whenever there is a change in the network (link failure, addition, new nodes, cost change or router cost out), the Type 1 and Type 10 LSA should be in sync. In most of the cases Type 1 and Type 10 are in sync, however it was observed that when IGP-LDP sync is not established, the Type 1 indicates the high cost while the Type 10 shows the normal cost. Use the **mpls traffic-eng ldp-sync-update** command to synchronize Label Distribution Protocol (LDP)-OSPF Type 1 and Type 10 link-state advertisement (LSA).

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to enable LDP-IGP OSPF Type 1 and Type 10 link-state advertisement (LSA) synchronization:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router ospf o1
RP/0/RSP0/CPU0:router(config-ospf)# mpls traffic-eng ldp-sync-update
```

mpls traffic-eng router-id (OSPF)

To specify that the traffic engineering router identifier for the node is the IP address associated with a given Open Shortest Path First (OSPF) interface, use the **mpls traffic-eng router-id** command in the appropriate configuration mode. To disable this feature, use the **no** form of this command.

mpls traffic-eng router-id {*router-id* | *type interface-path-id*}
no mpls traffic-eng router-id {*router-id* | *type interface-path-id*}

Syntax Description

router-id	The 32-bit router ID value specified in four-part, dotted-decimal notation (must be in the valid IP address range of 0.0.0.0 to 255.255.255.255).
type	Interface type. For more information, use the question mark (?) online help function.
interface-path-id	Physical interface or virtual interface. 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

If this command is specified in router configuration mode, then the traffic engineering router identifier for the node is the IP address associated with a given interface.

Command Modes

Router configuration

VRF 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

This identifier of the router acts as a stable IP address for the traffic engineering configuration. This IP address is flooded to all nodes. For all traffic engineering tunnels originating at other nodes and ending at this node, you must set the tunnel destination to the traffic engineering router identifier of the destination node, because that is the address that the traffic engineering topology database at the tunnel head uses for its path calculation.



Note We recommend that loopback interfaces be used for Multiprotocol Label Switching traffic engineering (MPLS TE), because they are more stable than physical interfaces.

Note This command is supported only in the default VRF mode.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to specify the traffic engineering router identifier as the IP address associated with loopback interface 0:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router ospf 1
RP/0/RSP0/CPU0:router(config-ospf)# mpls traffic-eng router-id loopback 0
```

Related Commands

Command	Description
mpls traffic-eng (OSPF), on page 1133	Configures an OSPF area for MPLS TE.

mtu-ignore (OSPF)

To prevent Open Shortest Path First (OSPF) from checking whether neighbors are using the same maximum transmission unit (MTU) on a common interface when exchanging database descriptor (DBD) packets, use the **mtu-ignore** command in the appropriate mode. To reset to default, use the **no** form of this command.

mtu-ignore [**disable** | **enable**]
no mtu-ignore

Syntax Description

disable	(Optional) Enables checking for whether OSPF neighbors are using the MTU on a common interface.
enable	(Optional) Disables checking for whether OSPF neighbors are using the MTU on a common interface.

Command Default

The default is **mtu-ignore** with no keywords, which disables MTU checking.

If this command is not specified in interface configuration mode,

then the interface adopts the MTU ignore parameter specified by the area.

If this command is not specified in area configuration mode,

then the interface adopts the MTU ignore parameter specified for the process.

If this command is not specified at any level,

then OSPF checks the MTU received from neighbors when exchanging DBD packets.

Command Modes

Interface configuration

Area configuration

Router configuration

VRF configuration

Multi-area 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

OSPF checks whether OSPF neighbors are using the same MTU on a common interface. This check is performed when neighbors exchange DBD packets. If the receiving MTU in the DBD packet is higher than the MTU configured on the incoming interface, OSPF adjacency is not established.

The keywords, **disable** and **enable** , do not need to be used. If no keywords are used, the **mtu-ignore** command disables MTU checking. You can then use the **no mtu-ignore** command to activate MTU checking.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to disable MTU mismatch detection on receiving DBD packets:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router ospf 109
RP/0/RSP0/CPU0:router(config-ospf)# area 0
RP/0/RSP0/CPU0:router(config-ospf-ar)# interface GigabitEthernet 0/1/0/3
RP/0/RSP0/CPU0:router(config-ospf-ar-if)# mtu-ignore
```

multi-area-interface

To enable multiple adjacencies for different Open Shortest Path First (OSPF) areas and enter multi-area interface configuration mode, use the **multi-area-interface** command in the area configuration mode. To reset to the default, use the **no** form of this command.

multi-area-interface *type interface-path-id*
no multi-area-interface *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

An OSPF network is enabled for one area only.

Command Modes

Area configuration

Command History

Release	Modification
Release 24.2.1	The command is extended to support loopback interfaces.
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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **multi-area-interface** command to enable area border routers (ABRs) to establish multiple adjacencies for different OSPF areas.

Each multiple area adjacency is announced as a point-to-point unnumbered link in the configured area. This point-to-point link provides a topological path for that area. The first or primary adjacency using the link advertises the link consistent with draft-ietf-ospf-multi-area-adj-06.txt.

You can configure multi-area adjacency on any interface where only two OSF speakers are attached. In the case of native broadcast networks, the interface must be configured as an OPSF point-to-point type using the **network point-to-point** command to enable the interface for a multi-area adjacency.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to enable multiple area adjacency for OSPF 109:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router ospf 109
RP/0/RSP0/CPU0:router(config-ospf)# area 0
RP/0/RSP0/CPU0:router(config-ospf-ar)# interface GigabitEthernet 0/1/0/3
RP/0/RSP0/CPU0:router(config-ospf-ar-if)# area 1
RP/0/RSP0/CPU0:router(config-ospf-ar)# multi-area-interface GigabitEthernet 0/1/0/3
RP/0/RSP0/CPU0:router(config-ospf-ar-mif)# ?

    authentication       Enable authentication
    authentication-key   Authentication password (key)
    commit               Commit the configuration changes to running
    cost                 Interface cost
    database-filter      Filter OSPF LSA during synchronization and flooding
    dead-interval        Interval after which a neighbor is declared dead
    describe             Describe a command without taking real actions
    distribute-list      Filter networks in routing updates
    do                   Run an exec command
    exit                 Exit from this submode
    hello-interval       Time between HELLO packets
    message-digest-key   Message digest authentication password (key)
    mtu-ignore           Enable/Disable ignoring of MTU in DBD packets
    no                   Negate a command or set its defaults
    packet-size          Customize size of OSPF packets upto MTU
    pwd                  Commands used to reach current submode
    retransmit-interval  Time between retransmitting lost link state advertisements
    root                 Exit to the global configuration mode
    show                 Show contents of configuration
    transmit-delay       Estimated time needed to send link-state update packet
RP/0/RSP0/CPU0:router(config-ospf-ar-mif)#
```

The following example shows how to configure, associate, and reuse a loopback IPv4 interface address for all areas in an ABR:

```
Router(config)#interface Loopback 0
Router(config-if)#ipv4 address 192.0.2.2/32
Router(config-if)#commit

Router(config)#router ospf 100
Router(config-ospf)#router-id 10.10.10.1
Router(config-ospf)#area 1
Router(config-ospf-ar)#interface Loopback 0
Router(config-ospf-ar-if)#commit

Router(config)#router ospf 100
Router(config-ospf)#router-id 10.10.10.1
Router(config-ospf)#area 0
Router(config-ospf-ar)#multi-area-interface Loopback 0
Router(config-ospf-ar-mif)#commit
```

Related Commands

Command	Description
show ospf interface, on page 1221	Displays OSPF interface information.

neighbor (OSPF)

To configure Open Shortest Path First (OSPF) routers interconnecting to nonbroadcast networks, use the **neighbor** command in interface configuration mode. To remove a configuration, use the **no** form of this command.

neighbor *ip-address* [**cost** *number*] [**priority** *number*] [**poll-interval** *seconds*]
no neighbor *ip-address* [**cost** *number*] [**priority** *number*] [**poll-interval** *seconds*]

Syntax Description

ip-address	Interface IP address of the neighbor.
cost *number*	(Optional) Assigns a cost to the neighbor, in the form of an integer from 1 to 65535. Neighbors with no specific cost configured assume the cost of the interface, based on the **cost** command. On point-to-multipoint interfaces, **cost** *number* is the only keyword and argument combination that works. The **cost** keyword does not apply to nonbroadcast multiaccess (NBMA) networks.
priority *number*	(Optional) Specifies an 8-bit number indicating the router priority value of the nonbroadcast neighbor associated with the IP address specified. The **priority** keyword does not apply to point-to-multipoint interfaces.
poll-interval *seconds*	(Optional) Specifies an unsigned integer value (in seconds) reflecting the poll interval. RFC 1247 recommends that this value be much larger than the hello interval. The **poll-interval** keyword does not apply to point-to-multipoint interfaces.

Command Default

No configuration is specified.

priority *number* : 0

poll-interval *seconds* : 120 seconds (2 minutes)

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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

You must include one neighbor entry in the software configuration for each known nonbroadcast network neighbor. The neighbor address must be on the primary address of the interface.

If a neighboring router has become inactive (hello packets have not been received for the router dead interval period), it may still be necessary to send hello packets to the dead neighbor. These hello packets are sent at a reduced rate called the *poll interval*.

When the router starts up, it sends only hello packets to those routers with nonzero priority; that is, routers that are eligible to become designated routers (DRs) and backup designated routers (BDRs). After the DR and BDR are selected, the DR and BDR start sending hello packets to all neighbors to form adjacencies.

To filter all outgoing OSPF link-state advertisement (LSA) packets for the neighbor, use the **neighbor database-filter all out** command.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to declare a router at address 172.16.3.4 on a nonbroadcast network, with a priority of 1 and a poll interval of 180 seconds:

```
  RP/0/RSP0/CPU0:router(config-ospf-ar-if)# neighbor 172.16.3.4 priority 1 poll-interval
180
```

The following example illustrates a network with nonbroadcast:

```
  RP/0/RSP0/CPU0:router# configure
  RP/0/RSP0/CPU0:router(config)# interface GigabitEthernet1/0/0/3
  RP/0/RSP0/CPU0:router(config-if)# ip address 172.16.3.10 255.255.255.0

  RP/0/RSP0/CPU0:router(config)# router ospf 1
  RP/0/RSP0/CPU0:router(config-ospf)# area 0
  RP/0/RSP0/CPU0:router(config-ospf-ar)# interface GigabitEthernet1/0/0/3
  RP/0/RSP0/CPU0:router(config-ospf-ar-if)# network nonbroadcast
  RP/0/RSP0/CPU0:router(config-ospf-ar-if)# neighbor 172.16.3.4 priority 1 poll-interval
180
  RP/0/RSP0/CPU0:router(config-ospf-ar-if)# neighbor 172.16.3.5 cost 10 priority 1
poll-interval 180
  RP/0/RSP0/CPU0:router(config-ospf-ar-if)# neighbor 172.16.3.6 cost 15 priority 1
poll-interval 180
  RP/0/RSP0/CPU0:router(config-ospf-ar-if)# neighbor 172.16.3.7 priority 1 poll-interval
180
```

Related Commands

Command	Description
neighbor database-filter all out, on page 1147	Filters all outgoing LSAs to an OSPF neighbor.
network (OSPF), on page 1148	Configures the OSPF network type to a type other than the default for a given medium.
priority (OSPF), on page 1166	Sets the router priority, which helps determine the designated router for this network.

neighbor database-filter all out

To filter all outgoing link-state advertisements (LSAs) to an Open Shortest Path First (OSPF) neighbor, use the **neighbor database-filter all out** command in interface configuration mode. To restore the forwarding of LSAs to the neighbor, use the **no** form of this command.

neighbor *ip-address* **database-filter all out**
no neighbor *ip-address* **database-filter all out**

Syntax Description

ip-address	IP address of the neighbor to which outgoing LSAs are blocked.

Command Default

Instead of all outgoing LSAs being filtered to the neighbor, they are flooded to the neighbor.

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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **neighbor database-filter all out** command to filter all outgoing OSPF LSA packets during synchronization and flooding for point-to-multipoint neighbors on nonbroadcast networks. More neighbor options are available with the **neighbor** command.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to prevent flooding of OSPF LSAs from point-to-multipoint networks to the neighbor at IP address 10.2.3.4:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router ospf 1
RP/0/RSP0/CPU0:router(config-ospf)# area 0
RP/0/RSP0/CPU0:router(config-ospf-ar)# interface GigabitEthernet1/0/0/3
RP/0/RSP0/CPU0:router(config-ospf-ar-if)# neighbor 10.2.3.4 database-filter all out
```

Related Commands

Command	Description
neighbor (OSPF), on page 1145	Configures OSPF routers interconnecting to nonbroadcast networks.

network (OSPF)

To configure the Open Shortest Path First (OSPF) network type to a type other than the default for a given medium, use the **network** command in the appropriate mode. To return to the default value, use the **no** form of this command.

network {**broadcast** | **non-broadcast** | {**point-to-multipoint** [**non-broadcast**] | **point-to-point**}}
no network

Syntax Description

broadcast	Sets the network type to broadcast.
non-broadcast	Sets the network type to nonbroadcast multiaccess (NBMA).
point-to-multipoint	Sets the network type to point-to-multipoint.
non-broadcast	(Optional) Sets the point-to-multipoint network to be nonbroadcast. If you use this keyword, the **neighbor** command is required.
point-to-point	Sets the network type to point-to-point.

Command Default

If this command is not specified in interface configuration mode, then the interface adopts the network parameter specified by the area.

If this command is not specified in area configuration mode, then the interface adopts the network parameter specified for the process.

If this command is not specified at any level, then the OSPF network type is the default of the given medium.

GigabitEthernt and TenGigEthernet interfaces default to broadcast.

Command Modes

Interface configuration

Area configuration

Router configuration

VRF 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **network** command to configure broadcast networks as NBMA networks when, for example, routers in your network do not support multicast addressing.

Configuring NBMA networks as either broadcast or nonbroadcast assumes that there are virtual circuits from every router to every router or fully meshed network. However, there are other configurations where this assumption is not true; for example, a partially meshed network. In these cases, you can configure the OSPF

network type as a point-to-multipoint network. Routing between two routers that are not directly connected go through the router that has virtual circuits to both routers.

If this command is issued on an interface that does not allow it, this command is ignored.

OSPF has two features related to point-to-multipoint networks. One feature applies to broadcast networks; the other feature applies to nonbroadcast networks:

- On point-to-multipoint, broadcast networks, you can use the **neighbor** command, and you must specify a cost to that neighbor.
- On point-to-multipoint, nonbroadcast networks, you must use the **neighbor** command to identify neighbors. Assigning a cost to a neighbor is optional.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to configure the OSPF network as a nonbroadcast network:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router ospf 1
RP/0/RSP0/CPU0:router(config-ospf)# area 0
RP/0/RSP0/CPU0:router(config-ospf-ar)# interface GigabitEthernet 0/1/0/3
RP/0/RSP0/CPU0:router(config-ospf-ar-if)# network non-broadcast
RP/0/RSP0/CPU0:router(config-ospf-ar-if)# neighbor 172.16.3.4 priority 1 poll-interval
180
```

Related Commands

Command	Description
neighbor (OSPF), on page 1145	Configures OSPF routers interconnecting to nonbroadcast networks.

nsf (OSPF)

To configure nonstop forwarding (NSF) for the Open Shortest Path First (OSPF) protocol, use the **nsf** command in the appropriate mode. To remove this command from the configuration file and restore the system to its default condition, use the **no** form of this command.

nsf {**cisco** [**enforce global**] | **ietf** [**helper disable**]}
no nsf {**cisco** [**enforce global**] | **ietf** [**helper disable**]}

Syntax Description

cisco	Enables Cisco Nonstop Forwarding.
enforce global	(Optional) Cancels NSF restart when non-NSF network device neighbors are detected.
ietf	Enables Internet Engineering Task Force (IETF) graceful restart.
helper disable	(Optional) Disables router helper support.

Command Default

NSF is disabled.

Command Modes

Router configuration

VRF 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

The NSF feature allows for the forwarding of data packets to continue along known routes while routing protocol information (such as OSPF) is being restored following a switchover.

Use the **nsf** command if the router is expected to perform NSF during restart. To experience the full benefits of this feature, configure all neighboring routers with NSF.

When this command is used without the optional **cisco enforce global** keywords and non-NSF neighbors are detected, the NSF restart mechanism aborts on the interfaces of those neighbors and functions properly on others.

When this command is used with the optional **cisco enforce global** keywords and non-NSF neighbors are detected, NSF restart is canceled for the entire OSPF process.

IETF graceful restart provides an NSF mechanism to allow data traffic to flow seamlessly with no packet drops during the transient period when OSPF attempts to recover after a process restart or RP failover, within the guidelines of RFC 3623.

By default, neighbors in helper mode listen to both the NSF Cisco- and NSF IETF-type LSAs. The **nsf** command enables one type of mechanism that would undergo an RP failover or, anticipating an OSPF process restart. If the **cisco** or **ietf** keyword is not entered, NSF is not enabled, irrespective of neighbors in listening mode for both NSF Cisco and NSF IETF.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to cancel NSF restart for the entire OSPF process if non-NSF neighbors are detected on any network interface during restart:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router ospf 1
RP/0/RSP0/CPU0:router(config-ospf)# nsf cisco enforce global
```

nsf flush-delay-time (OSPF)

To configure the maximum time allowed for nonstop forwarding (NSF) external route queries for the Open Shortest Path First (OSPF) protocol, use the **nsf flush-delay-time** command in the appropriate mode. To remove this command from the configuration file and restore the system to its default condition, use the **no** form of this command.

nsf flush-delay-time *seconds*
no nsf flush-delay-time *seconds*

Syntax Description

seconds	Length of time (in seconds) allowed for NSF external route queries. Range is 1 to 3600 seconds.

Command Default

seconds : 300

Command Modes

Router configuration

VRF 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 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
ospf	read, write

Examples

The following example shows how to configure the maximum time for NSF to learn external routes for OSPF at 60 seconds:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router ospf 1
RP/0/RSP0/CPU0:router(config-ospf)# nsf flush-delay-time 60
```

nsf interval (OSPF)

To configure the minimum time between consecutive nonstop forwarding (NSF) restart attempts for the Open Shortest Path First (OSPF) protocol, use the **nsf interval** command in the appropriate mode. To remove this command from the configuration file and restore the system to its default condition, use the **no** form of this command.

nsf interval *seconds*
no nsf interval *seconds*

Syntax Description

seconds	Length of time (in seconds) between consecutive restart attempts. Range is 90 to 3600 seconds.

Command Default

seconds : 90

Command Modes

Router configuration

VRF 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

When you use the **nsf interval** command, the OSPF process must be up for at least 90 seconds before OSPF attempts to perform an NSF restart.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to configure the minimum time between consecutive NSF restart attempts at 120 seconds:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:routerr(config)# router ospf 1
RP/0/RSP0/CPU0:router(config-ospf)# nsf interval 120
```

nsf lifetime (OSPF)

To configure the maximum time that routes are held in the Routing Information Base (RIB) following an Open Shortest Path First (OSPF) process restart, use the **nsf lifetime** command in the appropriate mode. To remove this command from the configuration file and restore the system to its default condition, use the **no** form of this command.

nsf lifetime *seconds*
no nsf lifetime *seconds*

Syntax Description

seconds	The length of time (in seconds) that routes are held in the RIB. Range is 90 to 3600 seconds.

Command Default

seconds : 95

Command Modes

Router configuration

VRF 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

When you use this command, the OSPF process must reconverge within the maximum length of time configured. If the convergence exceeds this length of time, routes are purged from RIB and nonstop forwarding (NSF) restart may fail.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to configure the maximum lifetime for OSPF NSF at 120 seconds:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router ospf 1
RP/0/RSP0/CPU0:router(config-ospf)# nsf lifetime 120
```

nsr (OSPF)

To configure nonstop routing (NSR) for the Open Shortest Path First (OSPF) protocol, use the **nsr** command in OSPF router configuration mode. To remove this command from the configuration file, use the **no** form of this command.

nsr [**disable**]
no nsr [**disable**]

Syntax Description

disable	(Optional) Disables NSR for all VRFs in this process.

Command Default

NSR is enabled.

Command Modes

Router configuration

Command History

Release	Modification
Release 3.7.2	This command was introduced.
Release 6.0.0	This command was modified. NSR was enabled by default.

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.

The NSR feature allows an OSPF process on the active RP to synchronize all necessary data and states with the OSPF process on the standby RP. When the switchover happens, the OSPF process on the newly active RP has all the necessary data and states to continue running and does not require any help from its neighbors.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to configure NSR:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router ospf 1
RP/0/RSP0/CPU0:router(config-ospf)# nsr
```

The following example shows how to disable NSR:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router ospf 1
RP/0/RSP0/CPU0:router(config-ospf)# nsr disable
```

nssa (OSPF)

To configure an area as a not-so-stubby area (NSSA), use the **nssa** command in area configuration mode. To remove the NSSA distinction from the area, use the **no** form of this command.

nssa [**no-redistribution**] [**default-information-originate** [**metric** *metric-value*] [**metric-type** *type-value*]] [**no-summary**]
no nssa

Syntax Description

no-redistribution	(Optional) Imports routes only into the normal areas, but not into the NSSA area, by the **redistribute** command when the router is an NSSA Area Border Router (ABR).
default-information-originate	(Optional) Generates a Type 7 default into the NSSA area. This keyword takes effect only on an NSSA ABR or NSSA Autonomous System Boundary Router (ASBR).
metric *metric-value*	(Optional) Specifies the metric used for generating the default route. If you omit a value and do not specify a value using the **defaultmetric** command, the default metric value is 10. Range is 1 to 16777214.
metric-type *type-value*	(Optional) Specifies the external link type associated with the default route advertised into the OSPF routing domain. It can be one of the following values: **1**—Type 1 external route **2**—Type 2 external route
no-summary	(Optional) Prevents an ABR from sending summary link advertisements into the NSSA.

Command Default

No NSSA area is defined.

Command Modes

Area 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

An NSSA does not flood Type 5 external LSAs from the core into the area, but can import autonomous system external routes in a limited fashion within the area.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to configure area 1 as an NSSA area:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router ospf 1
RP/0/RSP0/CPU0:router(config-ospf)# area 1
RP/0/RSP0/CPU0:router(config-ospf-ar)# nssa
```

ospf name-lookup

To configure the Open Shortest Path First (OSPF) protocol to look up Domain Name System (DNS) names, use the **ospf name-lookup** command in global configuration mode. To disable this function, use the **no** form of this command.

ospf name-lookup
no ospf name-lookup

Command Default

Routers are displayed by router ID or neighbor ID.

Command Modes

Global 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **ospf name-lookup** command to easily identify a router when executing all OSPF **show** command displays. The router is displayed by name rather than by its router ID or neighbor ID.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to configure OSPF to identify a router by name:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# ospf name-lookup
```

packet-size (OSPF)

To configure the size of Open Shortest Path First (OSPF) packets up to the size specified by the maximum transmission unit (MTU), use the **packet-size** command in the appropriate configuration mode. To disable this function and reestablish the default packet size, use the **no** form of this command.

packet-size *bytes*
no packet-size

Syntax Description

bytes	Size, in bytes. Range is 576 to 10000 bytes.

Command Default

If the command is not specified, the default packet size is either the interface IP MTU size (if that is lower than 9000 bytes) or 9000 bytes.

Command Modes

Router configuration

Area configuration

Interface configuration

VRF configuration

Multi-area configuration

Command History

Release	Modification
Release 3.7.2	This command was introduced.
Release 3.9.0	The default packet size was changed to the lower interface IP MTU size or 9000 bytes.

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.

Use the **packet-size** command to customize the size of OSPF packets. The OSPF protocol compares the packet size and the MTU size and uses the lower packet size value.

If the command is not configured, the default packet size is equal to the interface IP MTU size (if that is lower than 9000 bytes) or 9000 bytes. For example, if the interface IP MTU size is 1500 bytes, OSPF uses packet size of 1500 bytes on the interface because the byte size is lower than 9000 bytes. If the interface IP MTU size is 9500 bytes, OSPF uses packet size of 9000 bytes on the interface because the byte size exceeds 9000 bytes. The interface IP MTU size depends on the interface and the platform. In most cases, the default interface IP MTU value will be lower than 9000 bytes.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to configure the packet size on an interface:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router ospf 1
RP/0/RSP0/CPU0:router(config-ospf)# area 0
RP/0/RSP0/CPU0:router(config-ospf-ar)# interface GigabitEthernet 1/0/0/2
RP/0/RSP0/CPU0:router(config-ospf-ar-if)# packet-size 3500
```

passive (OSPF)

To suppress the sending of Open Shortest Path First (OSPF) protocol operation on an interface, use the **passive** command in the appropriate mode. To remove the passive configuration, use the **no** form of this command.

passive [**disable** | **enable**]
no passive

Syntax Description

disable	(Optional) Sends OSPF updates.
enable	(Optional) Disables sending OSPF updates.

Command Default

If this command is not specified in interface configuration mode, then the interface adopts the passive parameter specified by the area.

If this command is not specified in area configuration mode, then the interface adopts the passive parameter specified for the process.

If this command is not specified at any level, then the passive parameter is disabled and OSPF updates are sent on the interface.

Command Modes

Interface configuration

Area configuration

Router configuration

VRF 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

OSPF routing information is neither sent nor received through the specified interface. The interface appears as a stub network in the OSPF router (Type 1) link-state advertisement (LSA).

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows that GigabitEthernet interface 1/0/0/2 reduces OSPF updates because passive mode is enabled; however, GigabitEthernet interface 0/1/0/3 receives normal OSPF traffic flow:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router ospf 1
RP/0/RSP0/CPU0:router(config-ospf)# area 0
RP/0/RSP0/CPU0:router(config-ospf-ar)# interface GigabitEthernet1/0/0/2
RP/0/RSP0/CPU0:router(config-ospf-ar-if)# passive
RP/0/RSP0/CPU0:router(config-ospf-ar-if)# exit
RP/0/RSP0/CPU0:router(config-ospf-ar)# interface GigabitEthernet1/0/0/3
RP/0/RSP0/CPU0:router(config-ospf-ar-if)# end
```

prefix-suppression

To configure suppression of prefixes for an OSPF process, an OSPF area, or an OSPF interface, use the **prefix suppression** command.

To disable this feature, use the **no** prefix for this command, or configure the **disable** option.

prefix-suppression [**disable** | **secondary-address**]
no prefix-suppression

Syntax Description

(none)	Suppresses all prefixes.
disable	(Optional) Disables prefix suppression.
secondary-address	(Optional) Suppresses only secondary prefixes. **Note** Secondary prefix suppression is not applicable for OSPFv3.

Command Default

By default, prefix suppression is disabled.

Command Modes

router ospf configuration

router ospf area configuration

router ospf area interface configuration

Command History

Release	Modification
Release 6.1.1	This command was introduced.

Usage Guidelines

Transit-only networks that connect two routers are usually configured with routing IP addresses that are advertised in the Links State Advertisements (LSAs). However, these prefixes are not needed for data traffic. Suppressing these prefixes would reduce the number of links in LSAs, thereby improving convergence and also reducing the vulnerability of potential remote attacks.

Prefixes can be suppressed for an OSPF process, an OSPF area, or for specific interfaces of a router.

You can suppress prefixes by using the **prefix-suppression** command. When configured, the command applies to all prefixes configured at that level. You can configure this command for an entire router runing OSPF (ospf process), an OSPF area, or an OSPF interface.

Note

- Prefix suppression is not valid on loopback interfaces.
- Prefix suppression is supported on OSPF and OSPFv3 networks.

Task ID

Task ID	Operations
ospf	read, write

The following example shows how you can configure prefix suppression for an OSPF process.

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router ospf pfx
RP/0/RSP0/CPU0:router(config-ospf)# router-id 10.10.10.10
RP/0/RSP0/CPU0:router(config-ospf)# prefix-suppression
```

The following example shows how you can configure prefix suppression for an OSPFv3 process.

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router ospfv3 pfx
RP/0/RSP0/CPU0:router(config-ospfv3)# router-id 10.10.10.10
RP/0/RSP0/CPU0:router(config-ospfv3)# prefix-suppression
```

The following example shows how you can configure prefix suppression for an OSPF area.

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router ospf pfx
RP/0/RSP0/CPU0:router(config-ospf)# router-id 10.10.10.10
RP/0/RSP0/CPU0:router(config-ospf)# area 0
RP/0/RSP0/CPU0:router(config-ospf-ar)# prefix-suppression
```

The following example shows how you can configure prefix suppression for an OSPFv3 area.

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router ospfv3 pfx
RP/0/RSP0/CPU0:router(config-ospfv3)# router-id 10.10.10.10
RP/0/RSP0/CPU0:router(config-ospfv3)# area 0
RP/0/RSP0/CPU0:router(config-ospfv3-ar)# prefix-suppression
```

The following example shows how you can configure prefix suppression for an OSPF or OSPFv3 interface.

```
RP/0/RSP0/CPU0:router(config-ospf-ar)# interface GigabitEthernet 0/0/0/1
RP/0/RSP0/CPU0:router(config-ospf-ar-if)# network point-to-point
RP/0/RSP0/CPU0:router(config-ospf-ar-if)# prefix-suppression
```

The following example shows how you can disable prefix suppression for an OSPF or OSPFv3 interface.

```
RP/0/RSP0/CPU0:router(config-ospf-ar)# interface GigabitEthernet 0/0/0/1
RP/0/RSP0/CPU0:router(config-ospf-ar-if)# network point-to-point
RP/0/RSP0/CPU0:router(config-ospf-ar-if)# prefix-suppression disable
```

priority (OSPF)

To set the router priority for an interface, which helps determine the designated router for an Open Shortest Path First (OSPF) link, use the **priority** command in the appropriate mode. To return to the default value, use the **no** form of this command.

priority *value*
no priority *value*

Syntax Description

value	8-bit unsigned integer indicating the router priority value. Range is 0 to 255.

Command Default

If this command is not specified in interface configuration mode, then the interface adopts the priority parameter specified by the area.

If this command is not specified in area configuration mode, then the interface adopts the priority parameter specified for the process.

If this command is not specified at any level, then the default priority is 1.

Command Modes

Interface configuration

Area configuration

Router configuration

VRF 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

When two routers attached to a network both attempt to become the designated router, the one with the higher router priority takes precedence. If there is a tie, the router with the higher router ID takes precedence. A router with a router priority set to zero is ineligible to become the designated router or backup designated router. Router priority is configured only for interfaces to multiaccess networks (in other words, not point-to-point networks).

This priority value is used when you configure the Open Shortest Path First (OSPF) protocol for nonbroadcast networks using the **neighbor** command for OSPF.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows that priority is set through the **priority** and **neighbor** commands for Routers A and B and that the neighbor priority value must reflect that of the neighbor router:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# interface GigabitEthernet 0/1/0/1
RP/0/RSP0/CPU0:router(config-if)# ipv4 address 10.0.0.2 255.255.255.0
RP/0/RSP0/CPU0:router(config-if)# exit
RP/0/RSP0/CPU0:router(config)# router ospf 1
RP/0/RSP0/CPU0:router(config-ospf)# area 0
RP/0/RSP0/CPU0:router(config-ospf-ar)# interface GigabitEthernet 0/1/0/1
RP/0/RSP0/CPU0:router(config-ospf-ar-if)# network non-broadcast
RP/0/RSP0/CPU0:router(config-ospf-ar-if)# priority 4
RP/0/RSP0/CPU0:router(config-ospf-ar-if)# neighbor 10.0.0.1 priority 6
```

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# interface GigabitEthernet POS 0/2/0/1
RP/0/RSP0/CPU0:router(config-if)# ipv4 address 10.0.0.1 255.255.255.0
RP/0/RSP0/CPU0:router(config-if)# exit
RP/0/RSP0/CPU0:router(config)# router ospf 1
RP/0/RSP0/CPU0:router(config-ospf)# area 0
RP/0/RSP0/CPU0:router(config-ospf-ar)# interface GigabitEthernet 0/2/0/1
RP/0/RSP0/CPU0:router(config-ospf-ar-if)# network non-broadcast
RP/0/RSP0/CPU0:router(config-ospf-ar-if)# priority 6
RP/0/RSP0/CPU0:router(config-ospf-ar-if)# neighbor 10.0.0.2 priority 4
```

Related Commands

Command	Description
neighbor (OSPF), on page 1145	Configures OSPF routers interconnecting to nonbroadcast networks.
network (OSPF), on page 1148	Configures the OSPF network type to a type other than the default for a given medium.

protocol shutdown

To disable an instance of the Open Shortest Path First (OSPF) protocol so that it cannot form an adjacency on any interface, use the **protocol shutdown** command in the router configuration mode. To reenable the OSPF protocol, use the **no** form of this command.

protocol shutdown
no protocol shutdown

Command Default

No default behavior or values

Command Modes

Router 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **protocol shutdown** command to disable the OSPF protocol for a specific routing instance without removing any existing OSPF configuration parameters.

The OSPF protocol continues to run on the router and you can use the current OSPF configuration, but OSPF does not form any adjacencies on any interface.

This command is similar to performing the **no router ospf** command.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to disable the OSPF 1 instance:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router ospfv3 1
RP/0/RSP0/CPU0:router(config-ospf)# protocol shutdown
```

queue dispatch flush-lsa

To change the number of LSAs scheduled (rate-limited) for flushing, that are processed in each iteration, use the **queue dispatch flush-lsa** command in router configuration mode. To return to the system default value, use the **no** form of this command.

queue dispatch flush-lsa *count*
no queue dispatch flush-lsa

Syntax Description

count	Maximum number of LSAs flushed per run. Range is 30 to 3000.

Command Default

The default LSAs flushed per run is 150 (when the count is not configured).

Command Modes

Router configuration

Command History

Release	Modification
Release 3.9.0	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
ospf	read, write

Examples

The following example shows how to limit the number of LSAs flushed per run to 30:

```
RP/0/RSP0/CPU0:router(config-ospf)# queue dispatch flush-lsa 30
```

Use the show ospf message-queue, on page 1229 command to see the queue dispatch values, peak lengths, and limits.

Related Commands

Command	Description
queue dispatch incoming, on page 1171	Limits the number of continuous incoming events processed.
queue dispatch rate-limited-lsa, on page 1173	Sets the maximum number of rate-limited link-state advertisements (LSAs) processed per run.

Command	Description
queue dispatch spf-lsa-limit, on page 1175	Limits the number of summary or external Type 3 to Type 7 link-state advertisements (LSAs) processed per shortest path first (SPF) run.
queue limit, on page 1176	Sets the high watermark for incoming priority events.
show ospf message-queue, on page 1229	Displays the information about the queue dispatch values, peak lengths, and limits.

queue dispatch incoming

To limit the number of incoming packets (LSAUpdates, LSAcks, DBDs, LSRequests, and Hellos that trigger a change state) processed, use the **queue dispatch incoming** command in router configuration mode. To return to the system default value, use the **no** form of this command.

queue dispatch incoming *count*
no queue dispatch incoming

Syntax Description

count	Maximum number of continuous events processed. Range is 30 to 3000.

Command Default

The default incoming count is 300 packets (when the count is not configured).

Command Modes

Router configuration

Command History

Release	Modification
Release 3.9.0	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
ospf	read, write

Examples

The following example shows how limit the number of incoming packets processed to 500:

```
RP/0/RSP0/CPU0:router(config-ospf)# queue dispatch incoming 500
```

Use the show ospf message-queue, on page 1229 command to see the queue dispatch values, peak lengths, and limits.

Related Commands

Command	Description
queue dispatch rate-limited-lsa, on page 1173	Sets the maximum number of rate-limited link-state advertisements (LSAs) processed per run.
queue dispatch spf-lsa-limit, on page 1175	Limits the number of summary or external Type 3 to Type 7 link-state advertisements (LSAs) processed per shortest path first (SPF) run.

Command	Description
queue limit, on page 1176	Sets the high watermark for incoming priority events.
show ospf message-queue, on page 1229	Displays the information about the queue dispatch values, peak lengths, and limits.

queue dispatch rate-limited-lsa

To set the maximum number of rate-limited link-state advertisement (LSA) (re-)originations processed per run, use the **queue dispatch rate-limited-lsa** command in router configuration mode. To return to the system default value, use the **no** form of this command.

queue dispatch rate-limited-lsa *count*
no queue dispatch rate-limited-lsa

Syntax Description

count	Maximum number of rate-limited LSAs processed per run. Range is 30 to 3000.

Command Default

The default number of rate-limited LSAs processed per run is 300 (when this count is not configured).

Command Modes

Router configuration

Command History

Release	Modification
Release 3.9.0	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
ospf	read, write

Examples

The following example shows how to set the maximum number of rate-limited LSA (re-)originations processed per run to 300:

```
RP/0/RSP0/CPU0:router(config-ospf)# queue dispatch rate-limited-lsa 300
```

Related Commands

Command	Description
queue dispatch incoming, on page 1171	Limits the number of continuous incoming events processed.
queue dispatch spf-lsa-limit, on page 1175	Limits the number of summary or external Type 3 to Type 7 link-state advertisements (LSAs) processed per shortest path first (SPF) run.
queue limit, on page 1176	Sets the high watermark for incoming priority events.

Command	Description
show ospf message-queue, on page 1229	Displays the information about the queue dispatch values, peak lengths, and limits.

queue dispatch spf-lsa-limit

To change the maximum number of Type 3-4 and Type 5-7 link-state advertisements (LSAs) processed per shortest path first (SPF) iteration within a single SPF run, use the **queue dispatch spf-lsa-limit** command in router configuration mode. To return to the system default value, use the **no** form of this command.

queue dispatch spf-lsa-limit *count*
no queue dispatch spf-lsa-limit

Syntax Description

count	Maximum number of continuous Type 3-4 and Type 5-7 LSAs processed per SPF in each scheduled iteration within a single SPF run. Range is 30 to 3000.

Command Default

The default number of Type 3-4 and Type 5-7 processed per run is 150 LSAs (when this command is not configured).

Command Modes

Router configuration

Command History

Release	Modification
Release 3.9.0	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
ospf	read, write

Examples

The following example shows how to limit the number of continuous Type 3-4 and Type 5-7 LSAs processed by SPF per scheduling run, to 100:

```
RP/0/RSP0/CPU0:router(config-ospf)# queue dispatch spf-lsa-limit 100
```

Related Commands

Command	Description
queue dispatch incoming, on page 1171	Limits the number of continuous incoming events processed.
queue dispatch rate-limited-lsa, on page 1173	Sets the maximum number of rate-limited link-state advertisements (LSAs) processed per run
queue limit, on page 1176	Sets the high watermark for incoming priority events.
show ospf message-queue, on page 1229	Displays the information about the queue dispatch values, peak lengths, and limits.

queue limit

To set the high watermark for incoming events by priority, use the **queue limit** in router configuration mode. To return to the system default values, use the **no** form of this command.

queue limit {**high** | **medium** | **low**} *count*
no queue limit {**high** | **medium** | **low**}

Syntax Description

high	High watermark for incoming high-priority events (state-changing Hellos).
medium	High watermark for incoming medium-priority events (LSA ACK).
low	High watermark for incoming low-priority events (DBD/LSUpd/LSReq).
count	Maximum number of events per queue. Events are dropped when the priority queue size exceeds this value. Range is 1000 to 30000.

Command Default

High watermark: 9500 (when the corresponding configuration is not present).

Medium watermark: 9000 (when the corresponding configuration is not present).

Low watermark: 8000 (when the corresponding configuration is not present).

Command Modes

Router configuration

Command History

Release	Modification
Release 3.9.0	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.

Always keep the limits in the following order of priority:

Limit for High > Limit for Medium > Limit for Low

Task ID

Task ID	Operations
ospf	read, write

Examples

The following examples show how to set the maximum number of events per queue:

```
RP/0/RSP0/CPU0:router(config-ospf)# queue limit high 11000
RP/0/RSP0/CPU0:router(config-ospf)# queue limit medium 10000
RP/0/RSP0/CPU0:router(config-ospf)# queue limit low 9000
```

Related Commands

Command	Description
queue dispatch incoming, on page 1171	Limits the number of continuous incoming events processed.
queue dispatch rate-limited-lsa, on page 1173	Sets the maximum number of rate-limited link-state advertisements (LSAs) processed per run.
queue dispatch spf-lsa-limit, on page 1175	Limits the number of summary or external Type 3 to Type 7 link-state advertisements (LSAs) processed per shortest path first (SPF) run.
show ospf message-queue, on page 1229	Displays the information about the queue dispatch values, peak lengths, and limits.

range (OSPF)

To consolidate and summarize routes at an area boundary, use the **range** command in area configuration mode. To disable this function, use the **no** form of this command.

range *ip-address mask* [**advertise** | **not-advertise**]
no range *ip-address mask* [**advertise** | **not-advertise**]

Syntax Description

ip-address	IP address in four-part, dotted-decimal notation.
mask	IP address mask.
advertise	(Optional) Sets the address range status to advertise and generates a Type 3 summary link-state advertisement (LSA).
not-advertise	(Optional) Sets the address range status to DoNotAdvertise. The Type 3 summary LSA is suppressed and the component networks remain hidden from other networks.

Command Default

When this command is not specified for Area Border Routers (ABRs), routes at an area boundary are not consolidated or summarized.

Advertise is the default.

Command Modes

Area 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **range** command only with Area Border Router (ABRs). Use the command to consolidate or summarize routes for an area. The result is that a single summary route is advertised to other areas by the ABR. Routing information is condensed at area boundaries. External to the area, a single route is advertised for each address range. This process is called *route summarization* .

Multiple range configurations specifying the **range** command can be configured. Thus, the OSPF protocol can summarize addresses for many different sets of address ranges.

The summarized route uses the maximum cost of the routes assumed in the range.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows area 36.0.0.0 consisting of interfaces whose IP addresses have "10.31.x.x" as the first two octets. The **range** command summarizes interfaces. Instead of advertising eight networks individually, the single route 10.31.0.0 255.255.0.0 is advertised:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router ospf 201
RP/0/RSP0/CPU0:router(config-ospf)# area 0
RP/0/RSP0/CPU0:router(config-ospf-ar-if)# interface GigabitEthernet 0/3/0/2
!
RP/0/RSP0/CPU0:router(config-ospf)# area 36.0.0.0
RP/0/RSP0/CPU0:router(config-ospf-ar)# range 10.31.0.0 255.255.0.0
RP/0/RSP0/CPU0:router(config-ospf-ar)# interface GigabitEthernet 0/1/0/0
RP/0/RSP0/CPU0:router(config-ospf-ar-if)# interface GigabitEthernet0/1/0/0
RP/0/RSP0/CPU0:router(config-ospf-ar-if)# interface GigabitEthernet0/1/0/1
RP/0/RSP0/CPU0:router(config-ospf-ar-if)# interface GigabitEthernet0/1/0/2
RP/0/RSP0/CPU0:router(config-ospf-ar-if)# interface GigabitEthernet0/1/0/3
RP/0/RSP0/CPU0:router(config-ospf-ar-if)# interface GigabitEthernet0/2/0/0
RP/0/RSP0/CPU0:router(config-ospf-ar-if)# interface GigabitEthernet0/2/0/1
RP/0/RSP0/CPU0:router(config-ospf-ar-if)# interface GigabitEthernet0/2/0/2
RP/0/RSP0/CPU0:router(config-ospf-ar-if)# interface GigabitEthernet0/2/0/3
RP/0/RSP0/CPU0:router(config-ospf-ar-if)# end
```

Related Commands

Command	Description
summary-prefix (OSPF), on page 1269	Creates aggregate addresses for routes being redistributed from another routing protocol into the OSPF protocol.

redistribute (OSPF)

To redistribute routes from one routing domain into Open Shortest Path First (OSPF), use the **redistribute** command in the appropriate mode. To remove the **redistribute** command from the configuration file and restore the system to its default condition in which the software does not redistribute routes, use the **no** form of this command.

Border Gateway Protocol (BGP)
redistribute bgp *process-id* [**preserve-med**] [**metric** *metric-value*] [**metric-type** {**1** | **2**}] [**route-policy** *policy-name*] [**tag** *tag-value*]
no redistribute bgp *process-id* [**metric** *metric-value*] [**metric-type** {**1** | **2**}] [**route-policy** *policy-name*] [**tag** *tag-value*]

Local Interface Routes
redistribute connected [**instance** *instance-name*] [**instance IPCP**][**metric** *metric-value*] [**metric-type** {**1** | **2**}] [**route-policy** *policy-name*] [**tag** *tag-value*]
no redistribute connected [**instance** *instance-name*] [**metric** *metric-value*] [**metric-type** {**1** | **2**}] [**route-policy** *policy-name*] [**tag** *tag-value*]

Directed-attached gateway redundancy (DAGR)
redistribute dagr [**metric** *metric-value*] [**metric-type** {**1** | **2**}] [**route-policy** *policy-name*] [**tag** *tag-value*]
no redistribute dagr [**metric** *metric-value*] [**metric-type** {**1** | **2**}] [**route-policy** *policy-name*] [**tag** *tag-value*]

Enhanced Interior Gateway Routing Protocol (EIGRP)
redistribute eigrp *process-id* [**match** {**external** [**1** | **2**] | **internal**}] [**metric** *metric-value*] [**metric-type** {**1** | **2**}] [**route-policy** *policy-name*] [**tag** *tag-value*]
no redistribute eigrp *process-id* [**match** {**external** [**1** | **2**] | **internal**}] [**metric** *metric-value*] [**metric-type** {**1** | **2**}] [**route-policy** *policy-name*] [**tag** *tag-value*]

Intermediate System-to-Intermediate System (IS-IS)
redistribute isis *process-id* [**level-1** | **level-2** | **level-1-2**] [**metric** *metric-value*] [**metric-type** {**1** | **2**}] [**route-policy** *policy-name*] [**tag** *tag-value*]
no redistribute isis *process-id* [**level-1** | **level-2** | **level-1-2**] [**metric** *metric-value*] [**metric-type** {**1** | **2**}] [**route-policy** *policy-name*] [**tag** *tag-value*]

Open Shortest Path First (OSPF)
redistribute ospf *process-id* [**match** {**external** [**1** | **2**] | **internal** | **nssa-external** [**1** | **2**]}] [**metric** *metric-value*] [**metric-type** {**1** | **2**}] [**route-policy** *policy-name*] [**tag** *tag-value*]
no redistribute ospf *process-id* [**match** {**external** [**1** | **2**] | **internal** | **nssa-external** [**1** | **2**]}] [**metric** *metric-value*] [**metric-type** {**1** | **2**}] [**route-policy** *policy-name*] [**tag** *tag-value*]

Routing Information Protocol (RIP)
redistribute rip [**metric** *metric-value*] [**metric-type** {**1** | **2**}] [**route-policy** *policy-name*] [**tag** *tag-value*]
no redistribute rip [**metric** *metric-value*] [**metric-type** {**1** | **2**}] [**route-policy** *policy-name*] [**tag** *tag-value*]

IP Static Routes
redistribute static [**metric** *metric-value*] [**metric-type** {**1** | **2**}] [**route-policy** *policy-name*] [**tag** *tag-value*]
no redistribute static [**metric** *metric-value*] [**metric-type** {**1** | **2**}] [**route-policy** *policy-name*] [**tag** *tag-value*]

Lsa-type summary
redistribute *protocol* [**lsa-type**] [**summary**]

Syntax Description

bgp	Distributes routes from the BGP protocol.
process-id	For the **bgp** keyword, an autonomous system number has the following ranges: • 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. For the **isis** keyword, an IS-IS instance name from which routes are to be redistributed. The value takes the form of a string. A decimal number can be entered, but it is stored internally as a string. For the **ospf** keyword, an OSPF instance name from which routes are to be redistributed. The value takes the form of a string. A decimal number can be entered, but it is stored internally as a string.
preserve-med	(Optional) Preserves the Multi Exit Discriminator (MED) of BGP routes.
metric *metric-value*	(Optional) Specifies the metric used for the redistributed route. Range is 1 to 16777214. Use a value consistent with the source protocol.
metric-type { **1** | **2** }	(Optional) Specifies the external link type associated with the route advertised into the OSPF routing domain. It can be one of two values: • **1** —Type 1 external route • **2** —Type 2 external route
tag *tag-value*	(Optional) Specifies the value attached to each external route. This value is not used by the OSPF protocol itself, but is carried in the external LSAs. Range is 0 to 4294967295.
route-policy *policy-name*	(Optional) Specifies the identifier of a configured policy. A policy is used to filter the importation of routes from this source routing protocol to OSPF.
connected	Distributes routes that are established automatically by virtue of having enabled IP on an interface.
instance	Connected instance.
instance-name	Name of the connected instance.
instance IPCP	Distributes routes from IPCP protocols.
eigrp	Distributes routes from the EIGRP protocol.

isis	Distributes routes from the IS-IS protocol.
level-1	(Optional) Redistributes Level 1 routes into other IP routing protocols independently.
l **evel-1-2**	(Optional) Distributes both Level 1 and Level 2 routes into other IP routing protocols.
level-2	(Optional) Distributes Level 2 routes into other IP routing protocols independently.
ospf	Distributes routes from the OSPF protocol.
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. If no match is specified, the default is no filtering.
rip	Distributes routes from the RIP protocol.
static	Distributes IP static routes.
dagr	Distributes routes from the directed-attached gateway redundancy (DAGR).
lsa-type	LSA type for redistributed routes.

Command Default

Route redistribution is disabled.

metric *metric-value:* Default is 20 for routes from all protocols except BGP routes, for which the default is 1.

metric-type : Type 2 external route.

Command Modes

Router configuration

VRF configuration

Command History

Release	Modification
Release 3.7.2	This command was introduced.

Release	Modification
Release 3.9.0	Asplain format for 4-byte Autonomous system numbers notation was supported. Support was added for redistribution from directed-attached gateway redundancy (DAGR). The keyword **dagr** was added. The **instance** keywrod and *instance-name* argument were added for connected routes.
Release 7.9.1	The command was modified. The **lsa-type** keyword was added to redistribute the routes from particular *protocol* as Type-3 LSAs.

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.

Note When redistributing routes (into OSPF) 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.

Redistributed routing information should always be filtered by the **policy** *policy-name* keyword and argument. This filtering ensures that only those routes intended by the administrator are redistributed into OSPF.

For information about routing policies, see the *Routing Policy Commands on the Cisco ASR 9000 Series Router* module of *Routing Command Reference for Cisco ASR 9000 Series Routers*.

Whenever you use the **redistribute** or default-information originate (OSPF), on page 1073 command to redistribute routes into an OSPF routing domain, the router automatically becomes an ASBR. However, an ASBR does not, by default, generate a default route into the OSPF routing domain.

When routes are redistributed between OSPF processes, no OSPF metrics are preserved.

When routes are redistributed into OSPF and no metric is specified with the **metric** keyword, OSPF uses 20 as the default metric for routes from all protocols except BGP routes, which get a metric of 1.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to cause BGP routes to be redistributed into an OSPF domain:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router ospf 110
RP/0/RSP0/CPU0:router(config-ospf)# redistribute bgp 100
```

The following example shows how to redistribute the specified IS-IS process routes into an OSPF domain. The IS-IS routes are redistributed with a metric of 100.

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router ospf 109
RP/0/RSP0/CPU0:router(config-ospf)# redistribute isis 108 metric 100
```

In the following example, network 10.0.0.0 appears as an external link-state advertisement (LSA) in OSPF 1:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# interface GigabitEthernet 0/1/0/1
RP/0/RSP0/CPU0:router(config-if)# ip address 10.0.0.0 255.0.0.0
!
RP/0/RSP0/CPU0:router(config)# interface GigabitEthernet 0/2/0/2
RP/0/RSP0/CPU0:router(config)# ip address 10.99.0.0 255.0.0.0
!
RP/0/RSP0/CPU0:router(config)# router ospf 1
RP/0/RSP0/CPU0:router(config-ospf)# redistribute ospf 2
RP/0/RSP0/CPU0:router(config-ospf)# area 0
RP/0/RSP0/CPU0:router(config-ospf-ar)# interface GigabitEthernet 0/2/0/2
!
RP/0/RSP0/CPU0:router(config)# router ospf 2
RP/0/RSP0/CPU0:router(config-ospf)# area 0
RP/0/RSP0/CPU0:router(config-ospf-ar)# interface GigabitEthernet 0/1/0/1
```

The following example displays the configuration of the **redistribute***protocol* **lsa-type summary** command is configured to redistribute the routes from particular *protocol* as Type-3 LSAs:

```
Router# config
Router(config)#router ospf 1
Router(config-ospf)#redistribute static lsa-type summary
Router(config-ospf)#summary-in enable
Router(config-ospf)#comit
```

Related Commands

Command	Description
default-information originate (OSPF), on page 1073	Generates a default external route into an OSPF routing domain.

retransmit-interval (OSPF)

To specify the time between link-state advertisement (LSA) retransmissions for adjacencies belonging to the Open Shortest Path First (OSPF) interface, use the **retransmit-interval** command in the appropriate mode. To return to the default value, use the **no** form of this command.

retransmit-interval *seconds*
no retransmit-interval

Syntax Description

seconds	Time (in seconds) between retransmissions. It must be greater than the expected round-trip delay between any two routers on the attached network. Range is 1 to 65535 seconds.

Command Default

If this command is not specified in interface configuration mode, then the interface adopts the retransmit interval parameter specified by the area.

If this command is not specified in area configuration mode, then the interface adopts the retransmit interval parameter specified for the process.

If this command is not specified at any level, then the default retransmit interval is 5 seconds.

Command Modes

Interface configuration

Area configuration

Router configuration

Virtual-link configuration

VRF configuration

Multi-area configuration

Sham-link 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

When a router sends an LSA to its neighbor, it keeps the LSA until it receives the acknowledgment message. If the router receives no acknowledgment, it resends the LSA.

The setting of this parameter should be conservative, or needless retransmission results. The value should be larger for serial lines and virtual links.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to set the retransmit interval value to 8 seconds in interface configuration mode:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router ospf 201
RP/0/RSP0/CPU0:router(config-ospf)# area 0
RP/0/RSP0/CPU0:router(config-ospf-ar)# interface GigabitEthernet 0/2/0/1
RP/0/RSP0/CPU0:router(config-ospf-ar-if)# retransmit-interval 8
```

route-policy (OSPF)

To specify a routing policy to filter Type 3 link-state advertisements (LSA), use the **route-policy** command in area configuration mode. To disable the routing policy, use the **no** form of this command.

route-policy *route-policy-name* {**in** | **out**}
no route-policy *route-policy-name* {**in** | **out**}

Syntax Description

route-policy-name	Name of route policy.
in	Applies policy to inbound routes.
out	Applies policy to outbound routes.

Command Default

No policy is applied.

Command Modes

Area 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **route-policy** command to specify an OSPF routing policy for an inbound or outbound route. The policy can be used to filter routes or modify route attributes.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to specify an OSPF route policy for inbound routes in area 0:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router ospf 109
RP/0/RSP0/CPU0:router(config-ospf)# area 0
RP/0/RSP0/CPU0:router(config-ospf-area)# route-policy area0_in in
```

router-id (OSPF)

To configure a router ID for the Open Shortest Path First (OSPF) process, use the **router-id** command in the appropriate mode. To cause the software to use the default method of determining the router ID, use the **no** form of this command after clearing or restarting the OSPF process.

router-id *router-id*
no router-id *router-id*

Syntax Description

router-id	32-bit router ID value specified in four-part, dotted-decimal notation.

Command Default

If this command is not configured, the router ID is the highest IP version 4 (IPv4) address for an interface on the router, with any loopback interface taking precedence.

Command Modes

Router configuration

VRF 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

It is good practice to use the **router-id** command to explicitly specify a unique 32-bit numeric value for the router ID. This action ensures that OSPF can function regardless of the interface address configuration. Clear the OSPF process using the **clear ospf process** command or restart the OSPF process for the **no router-id** command to take effect.

OSPF attempts to obtain a router ID in the following ways (in order of preference):

1. By default, when the OSPF process initializes, it checks if there is a router-id in the checkpointing database.
2. The 32-bit numeric value specified by the OSPF **router-id** command in router configuration mode. (This value can be any 32-bit value. It is not restricted to the IPv4 addresses assigned to interfaces on this router, and need not be a routable IPv4 address.)
3. The ITAL selected router-id.
4. The primary IPv4 address of an interface over which this OSPF process is running. The first interface address in the OSPF interface is selected.

Note Unlike OSPF version 3, OSPF version 2 is guaranteed to have at least one interface with an IPv4 address configured.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to assign the IP address of 172.20.10.10 to the OSPF process 109:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router ospf 109
RP/0/RSP0/CPU0:router(config-ospf)# router-id 172.20.10.10
```

Related Commands

Command	Description
clear ospf process, on page 1057	Resets an OSPF router process without stopping and restarting it.
ipv4 address	Sets a primary IPv4 address for an interface.

router ospf (OSPF)

To configure an Open Shortest Path First (OSPF) routing process, use the **router ospf** command in global configuration mode. To terminate an OSPF routing process, use the **no** form of this command.

router ospf *process-name*
no router ospf *process-name*

Syntax Description

process-name	Name that uniquely identifies an OSPF routing process. The process name is any alphanumeric string no longer than 40 characters without spaces.

Command Default

No OSPF routing process is defined.

Command Modes

global 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

You can specify multiple OSPF routing processes in each router. Up to 10 processes can be configured. The recommendation is not to exceed 4 OSPF processes.

All OSPF configuration commands must be configured under an OSPF routing process. For example, two of these commands are the **default-metric** command and the **router-id** command.

Task ID

Task ID	Operations
ospf	read, write
rib	read, write

Examples

The following example shows how to instantiate an OSPF routing process called 109:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router ospf 109
```

Related Commands

Command	Description
area (OSPF), on page 1047	Configures an OSPF area.

Command	Description
default-metric (OSPF), on page 1075	Sets default metric values for routes redistributed from another protocol into the OSPF protocol.
interface (OSPF), on page 1108	Defines the interfaces on which the OSPF protocol runs.
router-id (OSPF), on page 1188	Configures a router ID for the OSPF process.

security ttl (OSPF)

To set the security time-to-live (TTL) value in the IP header for Open Shortest Path First (OSPF) packets, use the **security ttl** command in the appropriate configuration mode. To remove this command from the configuration file and restore the system to its default condition, use the **no** form of this command.

security ttl [**hops** *hops-number*]
no security ttl

Syntax Description

hops *hops-number*	IP hops. Maximum number of hops allowed. Range is 1 to 254 hops.

Command Default

hops-number : 1

Command Modes

Router configuration

Area configuration

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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

The **security ttl** command is used for the Generalized TTL Security Mechanism (GTSM) feature to prevent network attacks.

During the act of receiving Link State Advertisement (LSA) from neighbors, network attacks can occur because there are no checks that unicast or multicast packets are originating from a neighbor that is one hop away or multiple hops away over virtual links.

For virtual links, OSPF packets travel multiple hops across the network; hence, the TTL value can be decremented several times. For these type of links, a minimum TTL value must be allowed and accepted for multiple-hop packets.

To filter network attacks originating from invalid sources traveling over multiple hops, the GTSM, RFC 3682, is used to prevent the attacks. GTSM filters link-local addresses and allows for only one-hop neighbor adjacencies through the configuration of TTL value 255. The TTL value in the IP header is set to when OSPF packets are originated and checked on the received OSPF packets against the default GTSM TTL value 255 or the user configured GTSM TTL value, blocking unauthorized OSPF packets originated from TTL hops away.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to set the security TTL for an interface:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router ospf 1
RP/0/RSP0/CPU0:router(config-ospf)# area 0
RP/0/RSP0/CPU0:router(config-ospf-ar)# interface GigabitEthernet0/6/0/3
RP/0/RSP0/CPU0:router(config-ospf-ar-if)# security ttl 2
```

segment-routing sr-prefer prefix-list

To set the preference of segment routing (SR) labels over label distribution protocol (LDP) labels, use the **segment-routing sr-prefer prefix-list** command. The label preference is used to decide the labeled path that will be used in case both LDP and SR labeled paths are available. This only affects the traffic that comes as IP (un-labeled) or traffic that resolves over a labeled path (such as L3VPNs, P2VPNs and so on).



Note If ACL is used, OSPF signals preference of SR labels over LDP labels only for prefixes that match the ACL. If ACL is not used, SR labels preference is signaled for all prefixes.

segment-routing sr-prefer prefix-list *[acl-name]*

Syntax Description

prefix-list	Sets preference of SR labels over LDP labels.
[acl-name]	Name of access control list.

Command Default

LDP labels are preferred by default.

Command Modes

Router configuration

Area configuration

Command History

Release	Modification
Release 5.3.2	This command was introduced

Usage Guidelines

No specific guidelines impact the use of this command.

Task ID

Task ID	Operations
ospf	read, write

Examples

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#ipv4 prefix-list foo
RP/0/RSP0/CPU0:router(config)#router ospf 1
RP/0/RSP0/CPU0:router(config-ospf)# segment-routing sr-prefer prefix-list foo
  area 0
   interface Loopback0
    prefix-sid index 1
   !
   interface GigabitEthernet0/0/0/0
   !
   interface GigabitEthernet0/2/0/0
```

```
 !
 interface GigabitEthernet0/2/0/3
 !
!
area 1
 interface GigabitEthernet0/2/0/7
 !
```

sham-link

To configure an Open Shortest Path First OSPF sham link between two provider edge routers, use the **sham-link** command in VRF area configuration mode. To terminate an (OSPF) sham link, use the **no** form of this command.

sham-link *source-address destination-address*
no sham-link

Syntax Description

source-address	IP address of the local (source) sham-link endpoint specified in four-part, dotted-decimal notation.
destination-address	IP address of the remote (destination) sham-link endpoint specified in four-part, dotted-decimal notation.

Command Default

No sham link is configured.

Command Modes

VRF area 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **sham-link** command to configure a point-to-point connection between two provider edge (PE) routers creating an interconnect between two VPN sites (VPN backbone). Sham links are configured on PE provider edge (PE) routers in a Multiprotocol Label Switching (MPLS) VPN backbone.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to configure an OSPF sham link:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router ospf 109
RP/0/RSP0/CPU0:router(config_ospf)# vrf vrf_a
RP/0/RSP0/CPU0:router(config_ospf_vrf)# area 0
RP/0/RSP0/CPU0:router(config_ospf_vrf_ar)# sham-link 192.168.40.0 172.16.30.0
RP/0/RSP0/CPU0:router(config_ospf_vrf_ar_sl)# cost 23
```

Related Commands

Command	Description
area (OSPF), on page 1047	Configures an OSPF area.
cost (OSPF), on page 1064	Explicitly specifies the cost of the interface (network) for OSPF path calculation.
vrf (OSPF), on page 1282	Configures an OSPF VPN routing and forwarding (VRF) instance.

show lpts punt excessive-flow-trap ospf

To display the details of bad actor identified for ospf protocol, use the **show lpts punt excessive-flow-trap ospf** command in the Global Configuration mode.

show lpts punt excessive-flow-trap ospf

Command Default

None

Command Modes

Global Configuration mode

Command History

Release	Modification
Release 6.0.1	This command was introduced.

Usage Guidelines

You must be in a user group associated with a task group that includes the proper task IDs. The command reference guides include the task IDs required for each command. If you suspect user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Task ID

Task ID	Operations
lpts	read
basic-services	read-write

Examples

This example shows to the details of bad actor identified for ospf protocol:

```
RP/0/RSP0/CPU0:router# show lpts punt excessive-flow-trap ospf
Parent Interface: GigabitEthernet0/2/0/19          Src MAC Addr: 0000.6505.0102

         Intf Handle: 0x08000580                          Location: 0/2/CPU0
           Protocol: OSPF                                     Punt Reason: OSPF-mc-known

        Penalty Rate: 0 pps (all packets dropped)    Penalty Timeout: 15 mins

      Time Remaining: 10 mins 3 secs
```

Related Commands

Command	Description
show running-config lpts punt excessive-flow-trap, on page 537	Displays the running configuration for the Excessive Punt Flow Trap feature.

show ospf

To display general information about Open Shortest Path First (OSPF) routing processes, use the **show ospf** command in EXEC mode.

show ospf [*process-name*] [**vrf** {*vrf-name* | **all**}] [**summary**]

Syntax Description

process-name	(Optional) Name that uniquely identifies an OSPF routing process. The process name is defined by the **router ospf** command. If this argument is included, only information for the specified routing process is displayed.
vrf *vrf-name* **all**	(Optional) Specifies an OSPF VPN routing and forwarding (VRF) instance. The *vrf-name* argument can be specified as an arbitrary string. The strings "default" and "all" are reserved values of the *vrf-name* argument.
summary	(Optional) Displays OSPF summary information.

Command Default

IPv4 and unicast address prefixes

Command Modes

EXEC

Command History

Release	Modification
Release 3.7.2	This command was introduced.
Release 3.9.0	Asplain format for 4-byte Autonomous system numbers notation was supported. The input parameters and output were modified to display 4-byte autonomous system numbers and extended communities in either asplain or asdot notations.
	LSA type for redistributed routes example was added.

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.

Use the **show ospf** command to provide basic information about the OSPF processes running on the router. Additional options provide in-depth information.

Task ID

Task ID	Operations
ospf	read

Examples

The following is sample output from the **show ospf** command:

```
RP/0/RSP0/CPU0:router#show ospf

Routing Process "ospf 1" with ID 1.1.1.1
 Supports only single TOS(TOS0) routes
```

```
Supports opaque LSA
It is an area border router
Initial SPF schedule delay 5000 msecs
Minimum hold time between two consecutive SPFs 10000 msecs
Maximum wait time between two consecutive SPFs 10000 msecs
Initial LSA throttle delay 500 msecs
Minimum hold time for LSA throttle 5000 msecs
Maximum wait time for LSA throttle 5000 msecs
Minimum LSA interval 5000 msecs. Minimum LSA arrival 1 secs
Maximum number of configured interfaces 255
Number of external LSA 0. Checksum Sum 00000000
Number of opaque AS LSA 0. Checksum Sum 00000000
Number of DCbitless external and opaque AS LSA 0
Number of DoNotAge external and opaque AS LSA 0
Number of areas in this router is 2. 2 normal 0 stub 0 nssa
External flood list length 0
Non-Stop Forwarding enabled
   Area BACKBONE(0) (Inactive)
       Number of interfaces in this area is 2
       SPF algorithm executed 8 times
       Number of LSA 2. Checksum Sum 0x01ba83
       Number of opaque link LSA 0. Checksum Sum 00000000
       Number of DCbitless LSA 0
       Number of indication LSA 0
       Number of DoNotAge LSA 0
       Flood list length 0
   Area 1
       Number of interfaces in this area is 1
       SPF algorithm executed 9 times
       Number of LSA 2. Checksum Sum 0x0153ea
       Number of opaque link LSA 0. Checksum Sum 00000000
       Number of DCbitless LSA 0
       Number of indication LSA 0
       Number of DoNotAge LSA 0
       Flood list length 0
```

This table describes the significant fields shown in the display.

Table 88: show ospf Field Descriptions

Field	Description
Routing Process “ospf 201” with ID 172.22.110.200	OSPF process name.
Supports only	Number of types of service supported (Type 0 only).
It is	Types are internal, area border, or autonomous system boundary.
Redistributing External Routes from	Lists of redistributed routes, by protocol.
SPF schedule delay	Delay time of SPF calculations.
Minimum LSA interval	Minimum interval between LSAs.
Minimum LSA arrival	Minimum elapsed time between accepting an update for the same link-state advertisement (LSA).
external LSA	Total number of Type 5 LSAs in the LSDB.

Field	Description
opaque LSA	Total number of Type 10 LSAs in the LSDB.
DCbitless...AS LSA	Total number of Demand Circuit Type 5 and Type 11 LSAs.
DoNotAge...AS LSA	Total number of Type 5 and Type 11 LSAs with the DoNotAge bit set.
Number of areas	Number of areas in router, area addresses, and so on.
Area BACKBONE	Backbone is area 0.

The following example displays the LSA type for redistributed routes:

```
Router #show ospf
Thu Dec  8 18:16:48.332 IST

 Routing Process "ospf 1" with ID 192.168.0.1
 Role: Primary Active
 NSR (Non-stop routing) is Enabled
 Supports only single TOS(TOS0) routes
 Supports opaque LSA
 It is an autonomous system boundary router
 Maximum number of non self-generated LSA allowed 1000
    Current number of non self-generated LSA 804
    Threshold for warning message 60%
    Ignore-time 1 minutes, reset-time 2 minutes
    Ignore-count allowed 2, current ignore-count 0
 Redistributing External Routes from,
    static
 Maximum number of redistributed prefixes 1000
    Threshold for warning message 70%
    Current number of redistributed prefixes 100
```

show ospf border-routers

To display the internal Open Shortest Path First (OSPF) routing table entries to an Area Border Router (ABR) and Autonomous System Boundary Router (ASBR), use the **show ospf border-routers** command in EXEC mode.

show ospf [*process-name*] [**vrf** {*vrf-name* | **all**}] **border-routers** [*router-id*]

Syntax Description

process-name	(Optional) OSPF process name. If this argument is included, only information for the specified routing process is included.
vrf *vrf-name* **all**	(Optional) Specifies an OSPF VPN routing and forwarding (VRF) instance. The *vrf-name* argument can be specified as an arbitrary string. The strings "default" and "all" are reserved vrf-names.
router-id	(Optional) Router ID associated with the border router. The value of the *router-id* argument can be any 32-bit router ID value specified in four-part, dotted-decimal notation. No default exists.

Command Default

IPv4 and unicast address prefixes

Command Modes

EXEC

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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **show ospf border-routers** command to list all OSPF border routers visible to the specified processes and to ascertain the OSPF topology of the router.

Task ID

Task ID	Operations
ospf	read

Examples

The following is sample output from the **show ospf border-routers** command:

```
RP/0/RSP0/CPU0:router# show ospf border-routers

OSPF 1 Internal Routing Table

Codes: i - Intra-area route, I - Inter-area route

i 172.31.97.53 [1] via 172.16.1.53, GigabitEthernet 3/0/0/0, ABR/ASBR , Area 0, SPF 3
```

This table describes the significant fields shown in the display.

Table 89: show ospf border-routers Field Descriptions

Field	Description
i	Type of this route; i indicates an intra-area route, I an interarea route.
172.31.97.53	Router ID of destination.
[1]	Cost of using this route.
172.16.1.53	Next-Next hop toward the destination.
GigabitEthernet 3/0/0/0	Packets destined for 172.16.1.53 are sent over GigabitEthernet interface 3/0/0/0.
ABR/ASBR	Router type of the destination; it is either an Area Border Router (ABR) or Autonomous System Boundary Router (ASBR) or both.
Area 0	Area ID of the area from which this route is learned.
SPF 3	Internal number of the shortest path first (SPF) calculation that installs this route.

show ospf database

To display lists of information related to the Open Shortest Path First (OSPF) database for a specific router, use the **show ospf database** command in EXEC mode.

show ospf [*process-name*] [**vrf** { *vrf-name* | **all** }] [*area-id*] **database**
show ospf [*process-name*] [**vrf** { *vrf-name* | **all** }] [*area-id*] **database** [**adv-router** *ip-address*]
show ospf [*process-name*] [**vrf** { *vrf-name* | **all** }] [*area-id*] **database** [**asbr-summary**] [*link-state-id*]
show ospf [*process-name*] [**vrf** { *vrf-name* | **all** }] [*area-id*] **database** [**asbr-summary**] [*link-state-id*] [**internal**] [**adv-router** [*ip-address*]]
show ospf [*process-name*] [**vrf** { *vrf-name* | **all** }] [*area-id*] **database** [**asbr-summary**] [*link-state-id*] [**internal**] [**self-originate**]
show ospf [*process-name*] [**vrf** { *vrf-name* | **all** }] [*area-id*] **database** [**database-summary**]
show ospf [*process-name*] **database** [**database-summary**] [**detail**]
show ospf [*process-name*] **database** [**database-summary**] [**adv-router** *router ID*]
show ospf [*process-name*] [**vrf** { *vrf-name* | **all** }] [*area-id*] **database** [**external**] [*link-state-id*]
show ospf [*process-name*] [**vrf** { *vrf-name* | **all** }] [*area-id*] **database** [**external**] [*link-state-id*] [**internal**] [**adv-router** [*ip-address*]]
show ospf [*process-name*] [**vrf** { *vrf-name* | **all** }] [*area-id*] **database** [**external**] [*link-state-id*] [**internal**] [**self-originate**]
show ospf [*process-name*] [**vrf** { *vrf-name* | **all** }] [*area-id*] **database** [**network**] [*link-state-id*]
show ospf [*process-name*] [**vrf** { *vrf-name* | **all** }] [*area-id*] **database** [**network**] [*link-state-id*] [**internal**] [**adv-router** [*ip-address*]]
show ospf [*process-name*] [**vrf** { *vrf-name* | **all** }] [*area-id*] **database** [**network**] [*link-state-id*] [**internal**] [**self-originate**]
show ospf [*process-name*] [**vrf** { *vrf-name* | **all** }] [*area-id*] **database** [**nssa-external**] [*link-state-id*]
show ospf [*process-name*] [**vrf** { *vrf-name* | **all** }] [*area-id*] **database** [**nssa-external**] [*link-state-id*] [**internal**] [**adv-router** [*ip-address*]]
show ospf [*process-name*] [**vrf** { *vrf-name* | **all** }] [*area-id*] **database** [**nssa-external**] [*link-state-id*] [**internal**] [**self-originate**]
show ospf [*process-name*] [**vrf** { *vrf-name* | **all** }] [*area-id*] **database** [**opaque-area**] [*link-state-id*]
show ospf [*process-name*] [**vrf** { *vrf-name* | **all** }] [*area-id*] **database** [**opaque-area**] [*link-state-id*] [**internal**] [**adv-router**] [*ip-address*]
show ospf [*process-name*] [**vrf** { *vrf-name* | **all** }] [*area-id*] **database** [**opaque-area**] [*link-state-id*] [**internal**] [**self-originate**]
show ospf [*process-name*] [**vrf** { *vrf-name* | **all** }] [*area-id*] **database** [**opaque-as**] [*link-state-id*]
show ospf [*process-name*] [**vrf** { *vrf-name* | **all** }] [*area-id*] **database** [**opaque-as**] [*link-state-id*] [**internal**] [**adv-router** [*ip-address*]]
show ospf [*process-name*] [**vrf** { *vrf-name* | **all** }] [*area-id*] **database** [**opaque-as**] [*link-state-id*] [**internal**] [**self-originate**]
show ospf [*process-name*] [**vrf** { *vrf-name* | **all** }] [*area-id*] **database** [**opaque-link**] [*link-state-id*]

show ospf [*process-name*] [**vrf** { *vrf-name* | **all** }] [*area-id*] **database** [**opaque-link**] [*link-state-id*] [**internal**] [**adv-router** [*ip-address*]]
show ospf [*process-name*] [**vrf** { *vrf-name* | **all** }] [*area-id*] **database** [**opaque-link**] [*link-state-id*] [**internal**] [**self-originate**]
show ospf [*process-name*] [**vrf** { *vrf-name* | **all** }] [*area-id*] **database** [**router**] [*link-state-id*]
show ospf [*process-name*] [**vrf** { *vrf-name* | **all** }] [*area-id*] **database** [**router**] [**internal**] [**adv-router** [*ip-address*]]
show ospf [*process-name*] [**vrf** { *vrf-name* | **all** }] [*area-id*] **database** [**router**] [**internal**] [**self-originate**] [*link-state-id*]
show ospf [*process-name*] [**vrf** { *vrf-name* | **all** }] [*area-id*] **database** [**self-originate**]
show ospf [*process-name*] [**vrf** { *vrf-name* | **all** }] [*area-id*] **database** [**summary**] [*link-state-id*]
show ospf [*process-name*] [**vrf** { *vrf-name* | **all** }] [*area-id*] **database** [**summary**] [*link-state-id*] [**internal**] [**adv-router** [*ip-address*]]
show ospf [*process-name*] [**vrf** { *vrf-name* | **all** }] [*area-id*] **database** [**summary**] [*link-state-id*] [**internal**] [**self-originate**] [*link-state-id*]

Syntax Description

process-name	(Optional) OSPF process name that uniquely identifies an OSPF routing process. The process name is any alphanumeric string no longer than 40 characters. If this argument is included, only information for the specified routing process is included.
vrf	(Optional) Specifies an OSPF VPN routing and forwarding (VRF) instance.
vrf-name	(Optional) Name of the OSPF VRF. The *vrf-name* argument can be specified as an arbitrary string. The strings “default” and “all” are reserved VRF names.
all	(Optional) Specifies all OSPF VRF instances.
area-id	(Optional) Area number used to define the particular area.
adv-router *ip-address*	(Optional) Displays all LSAs of the specified router.
asbr-summary	(Optional) Displays information only about the Autonomous System Boundary Router (ASBR) summary LSAs.

link-state-id	(Optional) Portion of the Internet environment that is being described by the advertisement. The value entered depends on the link-state type of the advertisement. It must be entered in the form of an IP address. When the link-state advertisement (LSA) is describing a network, the *link-state-id* can take one of two forms: • The network IP address (as in Type 3 summary link advertisements and in autonomous system external link advertisements). • A derived address obtained from the link-state ID. **Note** Masking the link-state ID of a network link advertisement with the subnet mask of the network yields the IP address of the network. When the LSA is describing a router, the link-state ID is always the OSPF router ID of the described router. When an autonomous system external advertisement (LS Type = 5) is describing a default route, its link-state ID is set to Default Destination (0.0.0.0).
internal	(Optional) Displays internal LSA information.
self-originate	(Optional) Displays only self-originated LSAs (from the local router).
database-summary	(Optional) Displays how many of each type of LSA for each area there are in the database and the total.
external	(Optional) Displays information only about the external LSAs.
network	(Optional) Displays information only about the network LSAs.
nssa-external	(Optional) Displays information only about the not-so-stubby area (NSSA) external LSAs.
opaque-area	(Optional) Displays information about the opaque Type 10 LSAs. Type 10 denotes an area-local scope. Refer to RFC 2370 for more information on the opaque LSA options.
opaque-as	(Optional) Displays information about the opaque Type 11 LSAs. Type 11 denotes that the LSA is flooded throughout the autonomous system.
opaque-link	(Optional) Displays information about the opaque Type 9 LSAs. Type 9 denotes a link-local scope.
router	(Optional) Displays information only about the router LSAs.
summary	(Optional) Displays information only about the summary LSAs.
detail	(Optional) Displays information about the number of LSA counts per router.

Command Default IPv4 and unicast address prefixes

Command Modes EXEC

Command History

Release	Modification
Release 3.7.2	This command was introduced.
Release 5.3.0	show ospf database opaque-area command is extended to display extended Link LSA information.
Release 7.9.1	The **detail** keyword was added.

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.

The various forms of the **show ospf database** command deliver information about different OSPF link-state advertisements. This command can be used to examine the link-state database (LSD) and its contents. Each router participating in an area having identical database entries pertaining to that area (with the exception of LSAs that are being flooded). Numerous options (such as **network** and **router**) are used to display portions of the database.

Task ID

Task ID	Operations
ospf	read

Examples

The following is sample output from the **show ospf database** command when no arguments or keywords are used:

```
RP/0/RSP0/CPU0:router# show ospf database

OSPF Router with ID (172.20.1.11) (Process ID 1)

                Router Link States (Area 0)

 Link ID         ADV Router      Age         Seq#       Checksum Link count
 172.20.1.8       172.20.1.8       1381      0x8000010D    0xEF60   2
 172.20.1.11      172.20.1.11      1460      0x800002FE    0xEB3D   4
 172.20.1.12      172.20.1.12      2027      0x80000090    0x875D   3
 172.20.1.27      172.20.1.27      1323      0x800001D6    0x12CC   3

                Net Link States (Area 0)

 Link ID         ADV Router      Age         Seq#       Checksum
 172.22.1.27      172.20.1.27      1323      0x8000005B    0xA8EE
 172.22.1.11      172.20.1.11      1461      0x8000005B    0x7AC

                Type-10 Opaque Link Area Link States (Area 0)

 Link ID         ADV Router      Age         Seq#       Checksum Opaque ID
 10.0.0.0         172.20.1.11      1461      0x800002C8    0x8483     0
 10.0.0.0         172.20.1.12      2027      0x80000080    0xF858     0
 10.0.0.0         172.20.1.27      1323      0x800001BC    0x919B     0
 10.0.0.1         172.20.1.11      1461      0x8000005E    0x5B43     1
```

This table describes the significant fields shown in the display.

Table 90: show ospf database Field Descriptions

Field	Description
Link ID	Router ID number.
ADV Router	ID of the advertising router.
Age	Link-state age.
Seq#	Link-state sequence number (detects old or duplicate LSAs).
Checksum	Fletcher checksum of the complete contents of the LSA.
Link count	Number of interfaces detected for the router.
Opaque ID	Opaque LSA ID number.

The following is sample output from the **show ospf database** command with the **asbr-summary** keyword:

```
RP/0/RSP0/CPU0:router# show ospf database asbr-summary

OSPF Router with ID (192.168.0.1) (Process ID 300)

Summary ASB Link States (Area 0.0.0.0)

  LS age: 1463
  Options: (No TOS-capability)
  LS Type: Summary Links (AS Boundary Router)
  Link State ID: 172.17.245.1 (AS Boundary Router address)
  Advertising Router: 172.17.241.5
  LS Seq Number: 80000072
  Checksum: 0x3548
  Length: 28
  Network Mask: /0
 TOS: 0  Metric: 1
```

This table describes the significant fields shown in the display.

Table 91: show ospf database asbr-summary Field Descriptions

Field	Description
OSPF Router with ID	Router ID number.
Process ID	OSPF process name.
LS age	Link-state age.
Options	Type of service options (Type 0 only).
LS Type	Link-state type.

Field	Description
Link State ID	Link-state ID (ASBR).
Advertising Router	ID of the advertising router.
LS Seq Number	Link-state sequence (detects old or duplicate LSAs).
Checksum	Link-state checksum (Fletcher checksum of the complete contents of the LSA).
Length	Length (in bytes) of the LSAs.
Network Mask	Network mask implemented.
TOS	Type of service.
Metric	Link-state metric.

The following is sample output from the **show ospf database** command with the **external** keyword:

```
RP/0/RSP0/CPU0:router# show ospf database external

OSPF Router with ID (192.168.0.1) (Process ID 300)

                  Type-5 AS External Link States

  LS age: 280
  Options: (No TOS-capability)
  LS Type: AS External Link
  Link State ID: 172.17.0.0 (External Network Number)
  Advertising Router: 172.17.70.6
  LS Seq Number: 80000AFD
  Checksum: 0xC3A
  Length: 36
  Network Mask: 255.255.0.0
        Metric Type: 2 (Larger than any link state path)
        TOS: 0
        Metric: 1
        Forward Address: 0.0.0.0
        External Route Tag: 0
```

This table describes the significant fields shown in the display.

Table 92: show ospf database external Field Descriptions

Field	Description
OSPF Router with Router ID	Router ID number.
Process ID	OSPF process name.
LS age	Link-state age.
Options	Type of service options (Type 0 only).

Field	Description
LS Type	Link-state type.
Link State ID	Link-state ID (external network number).
Advertising Router	ID of the advertising router.
LS Seq Number	Link-state sequence number (detects old or duplicate LSAs).
Checksum	Link-state checksum (Fletcher checksum of the complete contents of the LSA).
Length	Length (in bytes) of the LSA.
Network Mask	Network mask implemented.
Metric Type	External type.
TOS	Type of service.
Metric	Link-state metric.
Forward Address	Forwarding address. Data traffic for the advertised destination is forwarded to this address. If the forwarding address is set to 0.0.0.0, data traffic is forwarded instead to the originator of the advertisement.
External Route Tag	External route tag, a 32-bit field attached to each external route. This tag is not used by the OSPF protocol itself.

The following is sample output from the **show ospf database** command with the **network** keyword:

```
RP/0/RSP0/CPU0:router# show ospf database network

 OSPF Router with ID (192.168.0.1) (Process ID 300)

Net Link States (Area 0.0.0.0)

  LS age: 1367
  Options: (No TOS-capability)
  LS Type: Network Links
  Link State ID: 172.23.1.3 (address of Designated Router)
  Advertising Router: 192.168.0.1
  LS Seq Number: 800000E7
  Checksum: 0x1229
  Length: 52
  Network Mask: /24
        Attached Router: 192.168.0.1
        Attached Router: 172.23.241.5
        Attached Router: 172.23.1.1
        Attached Router: 172.23.54.5
        Attached Router: 172.23.1.5
```

This table describes the significant fields shown in the display.

Table 93: show ospf database network Field Descriptions

Field	Description
OSPF Router with ID	Router ID number.
Process ID	OSPF process name.
LS age	Link-state age.
Options	Type of service options (Type 0 only).
LS Type	Link-state type.
Link State ID	Link-state ID of the designated router.
Advertising Router	ID of the advertising router.
LS Seq Number	Link-state sequence number (detects old or duplicate LSAs).
Checksum	Link-state checksum (Fletcher checksum of the complete contents of the LSA).
Length	Length (in bytes) of the LSA.
Network Mask	Network mask implemented.
Attached Router	List of routers attached to the network, by IP address.

The following is sample output, carrying Multiprotocol Label Switching traffic engineering (MPLS TE) specification information, from the **show ospf database** command with the **opaque-area** keyword and a *link-state-id* of adv-router:

```
RP/0/RSP0/CPU0:router# show ospf database opaque-area adv-router 172.20.1.12

  OSPF Router with ID (172.20.1.11) (Process ID 1)

                Type-10 Opaque Link Area Link States (Area 0)

   LS age: 224
   Options: (No TOS-capability, DC)
   LS Type: Opaque Area Link
   Link State ID: 1.0.0.0
   Opaque Type: 1
   Opaque ID: 0
   Advertising Router: 172.20.1.12
   LS Seq Number: 80000081
   Checksum: 0xF659
   Length: 132
   Fragment number : 0

     MPLS TE router ID : 172.20.1.12

     Link connected to Point-to-Point network
       Link ID : 172.20.1.11
       Interface Address : 172.21.1.12
       Neighbor Address : 172.21.1.11
       Admin Metric : 10
```

```
      Maximum bandwidth : 193000
      Maximum reservable bandwidth : 125000
      Number of Priority : 8
      Priority 0 : 125000      Priority 1 : 125000
      Priority 2 : 125000      Priority 3 : 125000
      Priority 4 : 125000      Priority 5 : 125000
      Priority 6 : 125000      Priority 7 : 100000
      Affinity Bit : 0x0

    Number of Links : 1
```

The following is sample output from the **show ospf database** command that displays a Type 10, Router Information LSA:

```
RP/0/RSP0/CPU0:router# show ospf database opaque-area 4.0.0.0

            OSPF Router with ID (3.3.3.3) (Process ID orange)

                Type-10 Opaque Link Area Link States (Area 0)

  LS age: 105
  Options: (No TOS-capability, DC)
  LS Type: Opaque Area Link
  Link State ID: 4.0.0.0
  Opaque Type: 4
  Opaque ID: 0
  Advertising Router: 3.3.3.3
  LS Seq Number: 80000052
  Checksum: 0x34e2
  Length: 52
  Fragment number: 0

    Router Information TLV: Length: 4
    Capabilities:
      Graceful Restart Helper Capable
      Traffic Engineering enabled area
      All capability bits: 0x50000000

    PCE Discovery TLV: Length: 20
      IPv4 Address: 3.3.3.3
      PCE Scope: 0x20000000
      Compute Capabilities:
      Inter-area default (Rd-bit)
      Compute Preferences:
      Intra-area: 0  Inter-area:  0
      Inter-AS:   0  Inter-layer: 0
```

This table describes the significant fields shown in the display.

Table 94: show ospf database opaque-area Field Descriptions

Field	Description
OSPF Router with ID	Router ID number.
Process ID	OSPF process name.
LS age	Link-state age.

Field	Description
Options	Type of service options (Type 0 only).
LS Type	Link-state type.
Link State ID	Link-state ID.
Opaque Type	Opaque link-state type.
Opaque ID	Opaque ID number.
Advertising Router	ID of the advertising router.
LS Seq Number	Link-state sequence (detects old or duplicate LSAs).
Checksum	Link-state checksum (Fletcher checksum of the complete contents of the LSA).
Length	Length (in bytes) of the LSA.
Fragment number	Arbitrary value used to maintain multiple traffic engineering LSAs.
Link ID	Link ID number.
Interface Address	ID address of the interface.
Neighbor Address	IP address of the neighbor.
Admin Metric	Administrative metric value used by MPLS TE.
Maximum bandwidth	Specifies maximum bandwidth (in kbps).
Maximum reservable bandwidth	Specifies maximum reservable bandwidth (in kbps).
Number of Priority	Priority number.
Affinity Bit	Used by MPLS TE.
Router Information TLV	Router capabilities are advertised in this TLV.
Capabilities	Some router capabilities include stub router, traffic engineering, graceful restart, and graceful restart helper.
PCE Discovery TLV	PCE address and capability information is advertised in this TLV.
IPv4 Address	Configured PCE IPv4 address.
PCE Scope	Computation capabilities of the PCE.
Compute Capabilities	Compute capabilities and preferences of the PCE.
Inter-area default (RD-bit)	PCE compute capabilities such as intra-area, inter-area, inter-area default, inter-AS, inter-AS default and inter-layer.

Field	Description
Compute Preferences	Order or preference of path computation that includes intra-area, inter-area, inter-AS, and inter-layer preferences.

The following is sample output from the **show ospf database** command with the **router** keyword:

```
RP/0/RSP0/CPU0:router# show ospf database router

OSPF Router with ID (192.168.0.1) (Process ID 300)

Router Link States (Area 0.0.0.0)

  LS age: 1176
  Options: (No TOS-capability)
  LS Type: Router Links
  Link State ID: 172.23.21.6
  Advertising Router: 172.23.21.6
  LS Seq Number: 80002CF6
  Checksum: 0x73B7
  Length: 120
  AS Boundary Router
  Number of Links: 8

 Link connected to: another Router (point-to-point)
(Link ID) Neighboring Router ID: 172.23.21.5
(Link Data) Router Interface address: 172.23.21.6
 Number of TOS metrics: 0
  TOS 0 Metrics: 2
```

This table describes the significant fields shown in the display.

Table 95: show ospf database router Field Descriptions

Field	Description
OSPF Router with ID	Router ID number.
Process ID	OSPF process name.
LS age	Link-state age.
Options	Type of service options (Type 0 only).
LS Type	Link-state type.
Link State ID	Link-state ID.
Advertising Router	ID of the advertising router.
LS Seq Number	Link-state sequence (detects old or duplicate LSAs).
Checksum	Link-state checksum (Fletcher checksum of the complete contents of the LSA).
Length	Length (in bytes) of the LSA.
AS Boundary Router	Definition of router type.

Field	Description
Number of Links	Number of active links.
Link ID	Link type.
Link Data	Router interface address.
TOS	Type of service metric (Type 0 only).

The following is sample output from **show ospf database** command with the **summary** keyword:

```
RP/0/RSP0/CPU0:router# show ospf database summary

       OSPF Router with ID (192.168.0.1) (Process ID 300)

 Summary Net Link States (Area 0.0.0.0)

  LS age: 1401
  Options: (No TOS-capability)
  LS Type: Summary Links (Network)
  Link State ID: 172.23.240.0 (Summary Network Number)
  Advertising Router: 172.23.241.5
  LS Seq Number: 80000072
  Checksum: 0x84FF
  Length: 28
  Network Mask: /24
     TOS: 0  Metric: 1
```

This table describes the significant fields shown in the display.

Table 96: show ospf database summary Field Descriptions

Field	Description
OSPF Router with ID	Router ID number.
Process ID	OSPF process name.
LS age	Link-state age.
Options	Type of service options (Type 0 only).
LS Type	Link-state type.
Link State ID	Link-state ID (summary network number).
Advertising Router	ID of the advertising router.
LS Seq Number	Link-state sequence (detects old or duplicate LSAs).
Checksum	Link-state checksum (Fletcher checksum of the complete contents of the LSA).
Length	Length (in bytes) of the LSA.
Network Mask	Network mask implemented.

Field	Description
TOS	Type of service.
Metric	Link-state metric.

The following is sample output from **show ospf database** command with the **database-summary** keyword:

```
RP/0/RSP0/CPU0:router# show ospf database database-summary

          OSPF Router with ID (172.19.65.21) (Process ID 1)

Area 0 database summary
  LSA Type      Count    Delete   Maxage
  Router        2        0        0
  Network       1        0        0
  Summary Net   2        0        0
  Summary ASBR  0        0        0
  Type-7 Ext    0        0        0
  Opaque Link   0        0        0
  Opaque Area   0        0        0
  Subtotal      5        0        0

Process 1 database summary
  LSA Type      Count    Delete   Maxage
  Router        2        0        0
  Network       1        0        0
  Summary Net   2        0        0
  Summary ASBR  0        0        0
  Type-7 Ext    0        0        0
  Opaque Link   0        0        0
  Opaque Area   0        0        0
  Type-5 Ext    2        0        0
  Opaque AS     0        0        0
  Total         7        0        0
```

This table describes the significant fields shown in the display.

Table 97: show ospf database database-summary Field Descriptions

Field	Description
LSA Type	Link-state type.
Count	Number of advertisements in that area for each link-state type.
Delete	Number of LSAs that are marked “Deleted” in that area.
Maxage	Number of LSAs that are marked “Maxaged” in that area.

The **show ospf database database-summary** with the **detail** keyword displays the number of LSA counts per router:

```
Router#show ospf database database-summary detail
Tue Dec  6 19:20:34.090 IST
```

```
            OSPF Router with ID (192.168.0.1) (Process ID 1)

Router 192.168.0.4 LSA summary
  LSA Type      Count    Delete   Maxage
  Router        0        0        0
  Network       0        0        0
  Summary Net   0        0        0
  Summary ASBR  0        0        0
  Type-5 Ext    697      0        0
  Type-7 Ext    0        0        0
  Opaque Link   0        0        0
  Opaque Area   0        0        0
  Opaque AS     6        0        0
  Total         703      0        0

Router 192.168.0.1 LSA summary
  LSA Type      Count    Delete   Maxage
  Router        1        0        0
  Network       0        0        0
  Summary Net   0        0        0
  Summary ASBR  0        0        0
  Type-5 Ext    0        0        0
  Type-7 Ext    0        0        0
  Opaque Link   0        0        0
  Opaque Area   64       0        0
  Opaque AS     0        0        0
  Total         65       0        0

Router 192.168.0.2 LSA summary
  LSA Type      Count    Delete   Maxage
  Router        1        0        0
  Network       0        0        0
  Summary Net   21       0        0
  Summary ASBR  2        0        0
  Type-5 Ext    0        0        0
  Type-7 Ext    0        0        0
  Opaque Link   0        0        0
  Opaque Area   21       0        0
  Opaque AS     0        0        0
  Total         45       0        0

Router 192.168.0.6 LSA summary
  LSA Type      Count    Delete   Maxage
  Router        1        0        0
  Network       0        0        0
  Summary Net   21       0        0
  Summary ASBR  2        0        0
  Type-5 Ext    0        0        0
  Type-7 Ext    0        0        0
  Opaque Link   0        0        0
  Opaque Area   19       0        0
  Opaque AS     0        0        0
  Total         43       0        0

Router 192.168.0.3 LSA summary
  LSA Type      Count    Delete   Maxage
  Router        0        0        0
  Network       0        0        0
  Summary Net   0        0        0
  Summary ASBR  0        0        0
  Type-5 Ext    7        0        0
  Type-7 Ext    0        0        0
```

```
  Opaque Link    0         0         0
  Opaque Area    0         0         0
  Opaque AS      6         0         0
  Total          13        0         0
```

The **show ospf database database-summary** with the **adv-router***router ID* keyword displays the router information and the LSAs received from a particular router:

```
Router#show ospf database database-summary adv-router 192.168.0.4
Tue Dec  6 19:21:04.549 IST


            OSPF Router with ID (192.168.0.1) (Process ID 1)

Router 192.168.0.4 LSA summary
  LSA Type       Count    Delete   Maxage
  Router         0        0        0
  Network        0        0        0
  Summary Net    0        0        0
  Summary ASBR   0        0        0
  Type-5 Ext     697      0        0
  Type-7 Ext     0        0        0
  Opaque Link    0        0        0
  Opaque Area    0        0        0
  Opaque AS      6        0        0
  Total          703      0        0
```

show ospf flood-list

To display a list of Open Shortest Path First (OSPF) link-state advertisements (LSAs) waiting to be flooded over an interface, use the **show ospf flood-list** command in EXEC mode.

show ospf [*process-name*] [**vrf** {*vrf-name* | **all**}] [*area-id*] **flood-list** [*type interface-path-id*]

Syntax Description

process-name	(Optional) OSPF process name that uniquely identifies an OSPF routing process. The process name is any alphanumeric string no longer than 40 characters. If this argument is included, only information for the specified routing process is included.
vrf	(Optional) Specifies an OSPF VPN routing and forwarding (VRF) instance.
vrf-name	(Optional) Name of the OSPF VRF. The *vrf-name* argument can be specified as an arbitrary string. The strings "default" and "all" are reserved VRF names.
all	(Optional) Specifies all OSPF VRF instances.
area-id	(Optional) Area number used to define the particular area.
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

All interfaces

Command Modes

EXEC

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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **show ospf flood-list** command to display LSAs in flood queue and queue length.

Flood list information is transient and normally the flood lists are empty.

Task ID

Task ID	Operations
ospf	read

Examples

The following is sample output from the **show ospf flood-list** command for interface GigabitEthernet 3/0/0/0:

```
RP/0/RSP0/CPU0:router# show ospf flood-list GigabitEthernet 3/0/0/0

 Interface GigabitEthernet3/0/0/0, Queue length 20
 Link state retransmission due in 12 msec
 Displaying 6 entries from flood list:

 Type  LS ID           ADV RTR         Seq NO      Age    Checksum
    5  10.2.195.0       200.0.0.163     0x80000009  0      0xFB61
    5  10.1.192.0       200.0.0.163     0x80000009  0      0x2938
    5  10.2.194.0       200.0.0.163     0x80000009  0      0x757
    5  10.1.193.0       200.0.0.163     0x80000009  0      0x1E42
    5  10.2.193.0       200.0.0.163     0x80000009  0      0x124D
    5  10.1.194.0       200.0.0.163     0x80000009  0      0x134C
```

This table describes the significant fields shown in the display.

Table 98: show ospf flood-list Field Descriptions

Field	Description
GigabitEthernet3/0/0/0	Interface for which information is displayed.
Queue length	Number of LSAs waiting to be flooded.
Link state retransmission due in	Length of time (in milliseconds) before next link-state transmission.
Type	Type of LSA.
LS ID	Link-state ID of the LSA.
ADV RTR	IP address of the advertising router.
Seq NO	Sequence number of the LSA.
Age	Age of the LSA (in seconds).
Checksum	Checksum of the LSA.

show ospf interface

To display Open Shortest Path First (OSPF) interface information, use the **show ospf interface** command in EXEC mode.

show ospf [*process-name*] [**vrf** {*vrf-name* | **all**}] [*area-id*] **interface** [**brief**] [*type interface-path-id*]

Syntax Description

process-name	(Optional) OSPF process name that uniquely identifies an OSPF routing process. The process name is any alphanumeric string no longer than 40 characters. If this argument is included, only information for the specified routing process is included.
vrf	(Optional) Specifies an OSPF VPN routing and forwarding (VRF) instance.
vrf-name	(Optional) Name of the OSPF VRF. The *vrf-name* argument can be specified as an arbitrary string. The strings "default" and "all" are reserved VRF names.
all	(Optional) Specifies all OSPF VRF instances.
area-id	(Optional) Area number used to define the particular area.
brief	(Optional) Displays brief interface information.
type	Interface type. For more information, use the question mark (?) online help function.
interface-path-id	Physical interface or virtual interface. 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

All interfaces

Command Modes

EXEC

Command History

Release	Modification
Release 3.7.2	This command was introduced.
Release 5.3.0	Command output extended to display the status of topology independent fast reroute (TI-FRR) on an interface.

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
ospf	read

Examples

The following is sample output from the **show ospf interface** command which includes the topology independent loop free alternates (TI LFA) related information:

```
RP/0/RSP0/CPU0:router# show ospf interface

 GigabitEthernet0/0/0/1 is up, line protocol is up
 Internet Address 1.2.2.1/24, Area 0
 Process ID 1, Router ID 0.0.0.1, Network Type BROADCAST, Cost: 1
 Transmit Delay is 1 sec, State BDR, Priority 1, MTU 1500, MaxPktSz 1500
 Designated Router (ID) 0.0.0.2, Interface address 1.2.2.2
 Backup Designated router (ID) 0.0.0.1, Interface address 1.2.2.
 Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
 Hello due in 00:00:02:857
 Index 2/2, flood queue length 0
 Next 0(0)/0(0)
 Last flood scan length is 1, maximum is 1
 Last flood scan time is 0 msec, maximum is 0 msec
 LS Ack List: current length 0, high water mark 6
 Neighbor Count is 1, Adjacent neighbor count is 1
 Adjacent with neighbor 0.0.0.2  (Designated Router)
 Suppress hello for 0 neighbor(s)
 Multi-area interface Count is 0
```

This table describes the significant fields shown in the display.

Table 99: show ospf interface Field Descriptions

Field	Description
GigabitEthernet	Status of the physical link.
line protocol	Operational status of the protocol.
Internet Address	Interface IP address, subnet mask, and area address.
Process ID	OSPF process ID, router ID, network type, and link-state cost.
Transmit Delay	Transmit delay, interface state, and router priority.
Timer intervals configured	Configuration of timer intervals.
Hello	Number of seconds until next hello packet is sent over this interface.
Index	Area and autonomous system flood indexes.
Next 0 (0) /0 (0)	Next area and autonomous system flood information, data pointer, and index.
Last flood scan length	Length of last flood scan.
Last flood scan time	Time (in milliseconds) of last flood scan.
Neighbor Count	Count of network neighbors and list of adjacent neighbors.
Suppress hello	Count of neighbors suppressing hello messages.

Field	Description
Multi-area interface	Multiple area interface information for the primary interface, such as count and area/neighbor locations.

show ospf mpls traffic-eng

To display information about the links and fragments available on the local router for traffic engineering, use the **show ospf mpls traffic-eng** command in EXEC mode.

show ospf [*process-name*] [**vrf** {*vrf-name* | **all**}] [*area-id*] [*type interface-path-id*] **mpls traffic-eng** {**link** | **fragment**}

Syntax Description

process-name	(Optional) OSPF process name that uniquely identifies an OSPF routing process. The process name is any alphanumeric string no longer than 40 characters. If this argument is included, only information for the specified routing process is included.
vrf *vrf-name* **all**	(Optional) Specifies an OSPF VPN routing and forwarding (VRF) instance. The *vrf-name* argument can be specified as an arbitrary string. The strings "default" and "all" are reserved VRF names.
area-id	(Optional) Area number used to define the particular area.
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.
link	Provides detailed information about the links over which traffic engineering is supported on the local router.
fragment	Provides detailed information about the traffic engineering fragments on the local router.

Command Default

All links or fragments

Command Modes

EXEC

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 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
ospf	read

Examples

The following is sample output from the **show ospf mpls traffic-eng** command when the **link** keyword is specified:

```
RP/0/RSP0/CPU0:router# show ospf mpls traffic-eng link

            OSPF Router with ID (10.10.10.10) (Process ID 1)

  Area 0 has 2 MPLS TE links. Area instance is 67441.

  Links in hash bucket 3.
    Link is associated with fragment 1. Link instance is 67441
      Link connected to Point-to-Point network
      Link ID : 10.10.10.8
      Interface Address : 10.10.10.2
      Neighbor Address : 10.10.10.3
      Admin Metric : 0
      Maximum bandwidth : 19440000
      Maximum global pool reservable bandwidth : 25000000
      Maximum sub pool reservable bandwidth    : 3125000
      Number of Priority : 8
      Global pool unreserved BW
      Priority 0 :   25000000  Priority 1 :   25000000
      Priority 2 :   25000000  Priority 3 :   25000000
      Priority 4 :   25000000  Priority 5 :   25000000
      Priority 6 :   25000000  Priority 7 :   25000000
      Sub pool unreserved BW
      Priority 0 :    3125000  Priority 1 :    3125000
      Priority 2 :    3125000  Priority 3 :    3125000
      Priority 4 :    3125000  Priority 5 :    3125000
      Priority 6 :    3125000  Priority 7 :    3125000
      Affinity Bit : 0


  Links in hash bucket 8.
    Link is associated with fragment 0. Link instance is 67441
      Link connected to Point-to-Point network
      Link ID : 10.1.1.1
      Interface Address : 10.10.25.4
      Neighbor Address : 10.10.25.5
      Admin Metric : 0
      Maximum bandwidth : 19440000
      Maximum global pool reservable bandwidth : 25000000
      Maximum sub pool reservable bandwidth    : 3125000
      Number of Priority : 8
      Global pool unreserved BW
      Priority 0 :   25000000  Priority 1 :   25000000
      Priority 2 :   25000000  Priority 3 :   25000000
      Priority 4 :   25000000  Priority 5 :   25000000
      Priority 6 :   25000000  Priority 7 :   25000000
      Sub pool unreserved BW
      Priority 0 :    3125000  Priority 1 :    3125000
      Priority 2 :    3125000  Priority 3 :    3125000
      Priority 4 :    3125000  Priority 5 :    3125000
      Priority 6 :    3125000  Priority 7 :    3125000
      Affinity Bit : 0
```

This table describes the significant fields shown in the display.

Table 100: show ospf mpls traffic-eng link Field Descriptions

Field	Description
Link ID	Link type.
Interface address	IP address of the interface.
Neighbor address	IP address of the neighbor.
Admin Metric	Administrative distance metric value used by Multiprotocol Label Switching traffic engineering (MPLS TE).
Maximum bandwidth	Bandwidth capacity of the link (in kbps).
Maximum global pool reservable bandwidth	Maximum amount of bandwidth that is available for reservation in the global pool.
Maximum sub pool reservable bandwidth	Maximum amount of bandwidth that is available for reservation in the subpool.
Number of Priority	Priority number.
Global pool unreserved BW	Amount of unreserved bandwidth that is available in the global pool.
Sub pool unreserved BW	Amount of unreserved bandwidth that is available in the subpool.
Affinity Bit	Used by MPLS TE. Attribute values required for links carrying this tunnel. A 32-bit dotted-decimal number. Valid values are from 0x0 to 0xFFFFFFFF, representing 32 attributes (bits), where the value of an attribute is 0 or 1.

The following is sample output from the **show ospf mpls traffic-eng** command when the **fragment** keyword is specified:

```
RP/0/RSP0/CPU0:router# show ospf mpls traffic-eng fragment

            OSPF Router with ID (10.10.10.10) (Process ID 1)

  Area 0 has 2 MPLS TE fragment. Area instance is 67441.
  MPLS router address is 10.10.10.10
  Next fragment ID is 2

  Fragment 0 has 1 link. Fragment instance is 67441.
  Fragment has 1 link the same as last update.
  Fragment advertise MPLS router address
    Link is associated with fragment 0. Link instance is 67441
      Link connected to Point-to-Point network
      Link ID : 10.1.1.1
      Interface Address : 10.10.25.4
      Neighbor Address : 10.10.25.5
      Admin Metric : 0
      Maximum bandwidth : 19440000
```

```
    Maximum global pool reservable bandwidth : 25000000
    Maximum sub pool reservable bandwidth    : 3125000
    Number of Priority : 8
    Global pool unreserved BW
    Priority 0 :   25000000  Priority 1 :   25000000
    Priority 2 :   25000000  Priority 3 :   25000000
    Priority 4 :   25000000  Priority 5 :   25000000
    Priority 6 :   25000000  Priority 7 :   25000000
    Sub pool unreserved BW
    Priority 0 :    3125000  Priority 1 :    3125000
    Priority 2 :    3125000  Priority 3 :    3125000
    Priority 4 :    3125000  Priority 5 :    3125000
    Priority 6 :    3125000  Priority 7 :    3125000
    Affinity Bit : 0


Fragment 1 has 1 link. Fragment instance is 67441.
Fragment has 0 link the same as last update.
  Link is associated with fragment 1. Link instance is 67441
    Link connected to Point-to-Point network
    Link ID : 10.10.10.8
    Interface Address : 10.10.10.2
    Neighbor Address : 10.10.10.3
    Admin Metric : 0
    Maximum bandwidth : 19440000
    Maximum global pool reservable bandwidth : 25000000
    Maximum sub pool reservable bandwidth    : 3125000
    Number of Priority : 8
    Global pool unreserved BW
    Priority 0 :   25000000  Priority 1 :   25000000
    Priority 2 :   25000000  Priority 3 :   25000000
    Priority 4 :   25000000  Priority 5 :   25000000
    Priority 6 :   25000000  Priority 7 :   25000000
    Sub pool unreserved BW
    Priority 0 :    3125000  Priority 1 :    3125000
    Priority 2 :    3125000  Priority 3 :    3125000
    Priority 4 :    3125000  Priority 5 :    3125000
    Priority 6 :    3125000  Priority 7 :    3125000
    Affinity Bit : 0
```

This table describes the significant fields shown in the display.

Table 101: show ospf mpls traffic-eng fragment Field Descriptions

Field	Description
Area instance	Number of times traffic engineering information or any link changed.
Link instance	Number of times any link changed.
Link ID	Link type.
Interface address	IP address of the interface.
Neighbor address	IP address of the neighbor.
Admin Metric	Administrative distance metric value used by MPLS TE.
Maximum bandwidth	Bandwidth capacity of the link (in kbps).

Field	Description
Maximum global pool reservable bandwidth	Maximum amount of bandwidth that is available for reservation in the global pool.
Maximum sub pool reservable bandwidth	Maximum amount of bandwidth that is available for reservation in the subpool.
Number of Priority	Priority number.
Global pool unreserved BW	Amount of unreserved bandwidth that is available in the global pool.
Sub pool unreserved BW	Amount of unreserved bandwidth that is available in the subpool.
Affinity Bit	Used by MPLS TE. Attribute values required for links carrying this tunnel. A 32-bit dotted-decimal number. Valid values are from 0x0 to 0xFFFFFFFF, representing 32 attributes (bits), where the value of an attribute is 0 or 1.

show ospf message-queue

To display the information about the queue dispatch values, peak lengths, and limits, use the **show ospf message-queue** command in EXEC mode.

show ospf message-queue

This command has no arguments or keywords.

Command Default

No default behavior or values

Command Modes

EXEC

Command History

Release	Modification
Release 3.9.0	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
ospf	read

Examples

The following is sample output from the **show ospf message-queue** command:

```
RP/0/RSP0/CPU0:router# show ospf 1 message-queue

OSPF 1
  Hello Input Queue:
    Current queue length: 0
    Event scheduled: 0
    Total queuing failures: 0
    Maximum length : 102
    Pkts pending processing: 0
    Limit: 5000

  Router Message Queue
    Current instance queue length: 0
    Current redistribution queue length: 0
    Current ex spf queue length: 0
    Current sum spf queue length: 0
    Current intra spf queue length: 0
    Event scheduled: 0
    Maximum length : 101
    Total low queuing failures: 0
    Total medium queuing failures: 0
    Total high queuing failures: 0
    Total instance events: 919
    Processing quantum : 300
    Low queuing limit: 8000
```

```
    Medium queuing limit: 9000
    High queuing limit: 9500
    Rate-limited LSA processing quantum: 150
    Current rate-limited LSA queue length: 0
    Rate-limited LSA queue peak len: 517

    Rate-limited LSAs processed: 4464
    Flush LSA processing quantum: 150
    Current flush LSA queue length: 0
    Flush LSA queue peak len: 274
    Rate-limited flush LSAs processed: 420

    SPF-LSA-limit processing quantum: 150
    Managed timers processing quantum: 50
    Instance message count: 0
    Instance pulse send count: 919
    Instance pulse received count: 919
    Global pulse count: 0
    Instance Pulse errors: 0

  TE Message Queue
    Current queue length: 0
    Total queuing failures: 0
    Maximum length : 0

Number of Dlink errors: 0
```

This table describes the significant fields shown in the display.

Table 102: show ospf message-queue Field Descriptions

Field	Description
Hello Input Queue	This section provides statistics on the number of events and incoming packets processed in the Hello (incoming packet) thread of the OSPF process.
Router Message Queue	This section provides statistics on the events and messages processed in the Router (primary) thread of the OSPF process.
TE Message Queue	This section provides statistics on traffic-engineering events and messages received by OSPF from TE (the te_control process). These events are processed in the Router thread of the OSPF process.
Number of Dlink errors	The number of enqueuing or dequeuing errors seen across all the linked-lists in the OSPF process.

Related Commands

Command	Description
queue dispatch incoming, on page 1171	Limits the number of continuous incoming events processed.
queue dispatch rate-limited-lsa, on page 1173	Sets the maximum number of rate-limited link-state advertisements (LSAs) processed per run.
queue dispatch spf-lsa-limit, on page 1175	Limits the number of summary or external Type 3 to Type 7 link-state advertisements (LSAs) processed per shortest path first (SPF) run.

Command	Description
queue limit, on page 1176	Sets the high watermark for incoming priority events.

show ospf neighbor

To display Open Shortest Path First (OSPF) neighbor information on an individual interface basis, use the **show ospf neighbor** command in EXEC mode.

show ospf [*process-name*] [**vrf** {*vrf-name* | **all**}] [*area-id*] **neighbor** [[*type interface-path-id*] [*neighbor-id*] [**detail**] | **area-sorted** | **capabilities**]

Syntax Description

process-name	(Optional) Name that uniquely identifies an OSPF routing process. The process name is defined by the **router ospf** command. If this argument is included, only information for the specified routing process is displayed.
vrf *vrf-name* **all**	(Optional) Specifies an OSPF VPN routing and forwarding (VRF) instance. The *vrf-name* argument can be specified as an arbitrary string. The strings "default" and "all" are reserved VRF names.
area-id	(Optional) Area ID. If you do not specify an area, all areas are displayed.
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.
neighbor-id	(Optional) Neighbor ID.
detail	(Optional) Displays all neighbors given in detail (lists all neighbors).
area-sorted	(Optional) Specifies that all neighbors are grouped by area.
capabilities	(Optional) Displays the capabilties of all neighbors advertised in the Router Information Link State Advertisement (RI LSA).

Command Default

All neighbors

Command Modes

EXEC

Command History

Release	Modification
Release 24.2.1	The **capabilities** keyword was introduced.
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 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
ospf	read

Examples

The following is sample output from the **show ospf neighbor** command showing two lines of summary information for each neighbor:

```
RP/0/RSP0/CPU0:router# show ospf neighbor

Neighbors for OSPF

Neighbor ID     Pri  State          Dead Time  Address           Interface
192.168.199.137 1    FULL/DR        0:00:31    172.31.80.37      GigabitEthernet 0/3/0/2
    Neighbor is up for 18:45:22
192.168.48.1     1    FULL/DROTHER  0:00:33    192.168.48.1       GigabitEthernet 0/3/0/3

    Neighbor is up for 18:45:30
192.168.48.200   1    FULL/DROTHER  0:00:33    192.168.48.200     GigabitEthernet 0/3/0/3

    Neighbor is up for 18:45:25
192.168.199.137 5    FULL/DR        0:00:33    192.168.48.189     GigabitEthernet 0/3/0/3
    Neighbor is up for 18:45:27
```

This table describes the significant fields shown in the display.

*Table 103: **show ospf neighbor** Field Descriptions*

Field	Description
Neighbor ID	Neighbor router ID.
Pri	Designated router priority.
State	OSPF state.
Dead time	Time (in hh:mm:ss) that must elapse before OSPF declares the neighbor dead.
Address	Address of next hop.
Interface	Interface name of next hop.
Neighbor is up	Amount of time (in hh:mm:ss) that the OSPF neighbor has been up.

The following is sample output showing summary information about the neighbor that matches the neighbor ID:

```
RP/0/RSP0/CPU0:router# show ospf neighbor 192.168.199.137
```

```
 Neighbor 192.168.199.137, interface address 172.31.80.37
    In the area 0.0.0.0 via interface GigabitEthernet 0/3/0/2
    Neighbor priority is 1, State is FULL, 6 state changes
    DR is 0.0.0.0 BDR is 0.0.0.0
    Options is 0x2
    Dead timer due in 0:00:32
    Neighbor is up for 18:45:30
    Number of DBD retrans during last exhange 0
    Index 1/1, retransmission queue length 0, number of retransmission 0
    First 0x0(0)/0x0(0) Next 0x0(0)/0x0(0)
    Last retransmission scan length is 0, maximum is 0
    Last retransmission scan time is 0 msec, maximum 0 msec
 Neighbor 192.168.199.137, interface address 192.168.48.189
    In the area 0.0.0.0 via interface GigabitEthernet 0/3/0/3
    Neighbor priority is 5, State is FULL, 6 state changes
    Options is 0x2
    Dead timer due in 0:00:32
    Neighbor is up for 18:45:30
    Number of DBD retrans during last exhange 0
    Index 1/1, retransmission queue length 0, number of retransmission 0
    First 0x0(0)/0x0(0) Next 0x0(0)/0x0(0)
    Last retransmission scan length is 0, maximum is 0
    Last retransmission scan time is 0 msec, maximum 0 msec

Total neighbor count: 2
```

This table describes the significant fields shown in the display.

Table 104: show ospf neighbor 192.168.199.137 Field Descriptions

Field	Description
Neighbor	Neighbor router ID.
interface address	IP address of the interface.
In the area	Area and interface through which the OSPF neighbor is known.
Neighbor priority	Router priority of neighbor and neighbor state.
State	OSPF state.
state changes	Number of state changes for this neighbor.
DR is	Neighbor ID of the designated router.
BDR is	Neighbor ID of the backup designated router.
Options	Hello packet options field contents(E-bit only; possible values are 0 and 2; 2 indicates area is not a stub; 0 indicates area is a stub.)
Dead timer	Time (in hh:mm:ss) to elapse before OSPF declares the neighbor dead.
Neighbor is up	Amount of time (in hh:mm:ss) that the OSPF neighbor has been up.
Number of DBD retrans	Number of re-sent database description packets.

Field	Description
Index	Index and the remaining lines of this command give detailed information about flooding information received from the neighbor.

If you specify the interface along with the neighbor ID, the software displays the neighbors that match the neighbor ID on the interface, as in the following sample display:

```
RP/0/RSP0/CPU0:router# show ospf neighbor GigabitEthernet 0/3/0/2 192.168.199.137

 Neighbor 192.168.199.137, interface address 172.31.80.37
    In the area 0.0.0.0 via interface GigabitEthernet 0/3/0/2
    Neighbor priority is 1, State is FULL, 6 state changes
    DR is 0.0.0.0 BDR is 0.0.0.0
    Options is 0x2
    Dead timer due in 0:00:32
    Neighbor is up for 18:45:30
    Number of DBD retrans during last exhange 0
    Index 1/1, retransmission queue length 0, number of retransmission 0
    First 0x0(0)/0x0(0) Next 0x0(0)/0x0(0)
    Last retransmission scan length is 0, maximum is 0
    Last retransmission scan time is 0 msec, maximum 0 msec

Total neighbor count: 1
```

This table describes the significant fields shown in the display.

Table 105: show ospf neighbor GigabitEthernet 0/3/0/2 192.168.199.137 Field Descriptions

Field	Description
Neighbor	Neighbor router ID.
interface address	IP address of the interface.
In the area	Area and interface through which the OSPF neighbor is known.
Neighbor priority	Router priority of the neighbor.
State	OSPF state.
state changes	Number of state changes for this neighbor.
DR is	Neighbor ID of the designated router.
BDR is	Neighbor ID of the backup designated router.
Options	Hello packet options field contents (E-bit only; possible values are 0 and 2; 2 indicates area is not a stub; 0 indicates area is a stub)
Dead timer	Time (in hh:mm:ss) to elapse before OSPF declares the neighbor dead.
Neighbor is up	Amount of time (in hh:mm:ss) that the OSPF neighbor has been up.
Number of DBD retrans	Number of re-sent database description packets.

Field	Description
Index	Index and the remaining lines of this command give detailed information about flooding information received from the neighbor.

You can also specify the interface without the neighbor ID to show all neighbors on the specified interface, as in the following sample display:

```
RP/0/RSP0/CPU0:router# show ospf neighbor GigabitEthernet POS 0/3/0/3

Neighbors for OSPF ospf1

    ID          Pri   State        Dead Time     Address         Interface
192.168.48.1     1    FULL/DROTHER  0:00:33    192.168.48.1       GigabitEthernet POS
0/3/0/3
     Neighbor is up for 18:50:52
192.168.48.200   1    FULL/DROTHER  0:00:32    192.168.48.200     GigabitEthernet POS
0/3/0/3
     Neighbor is up for 18:50:52
192.168.199.137 5    FULL/DR       0:00:32    192.168.48.189     GigabitEthernet POS
0/3/0/3
     Neighbor is up for 18:50:52

Total neighbor count: 3
```

This table describes the significant fields shown in the display.

Table 106: show ospf neighbor GigabitEthernet 0/3/0/3 Field Descriptions

Field	Description
ID	Neighbor router ID.
Pri	Route priority of the neighbor.
State	OSPF state.
Dead Time	Time (in hh:mm:ss) to elapse before OSPF declares the neighbor dead.
Address	Address of next hop.
Interface	Interface name of next hop.
Neighbor is up	Time (in hh:mm:ss) that the OSPF neighbor has been up.
Options	Hello packet options field contents (E-bit only; possible values are 0 and 2; 2 indicates area is not a stub; 0 indicates area is a stub)
Dead timer	Time (in hh:mm:ss) to elapse before OSPF declares the neighbor dead.
Neighbor is up	Amount of time (in hh:mm:ss) that the OSPF neighbor has been up.
Number of DBD retrans	Number of re-sent database description packets.
Index	Index and the remaining lines of this command give detailed information about flooding information received from the neighbor.

The following samples are from output from the **show ospf neighbor detail** command:

```
RP/0/RSP0/CPU0:router# show ospf neighbor detail

 Neighbor 192.168.199.137, interface address 172.31.80.37
    In the area 0.0.0.0 via interface GigabitEthernet 0/3/0/2
    Neighbor priority is 1, State is FULL, 6 state changes
    DR is 0.0.0.0 BDR is 0.0.0.0
    Options is 0x2
    Dead timer due in 0:00:32
    Neighbor is up for 18:45:30
    Number of DBD retrans during last exhange 0
    Index 1/1, retransmission queue length 0, number of retransmission 0
    First 0x0(0)/0x0(0) Next 0x0(0)/0x0(0)
    Last retransmission scan length is 0, maximum is 0
    Last retransmission scan time is 0 msec, maximum 0 msec

Total neighbor count: 1


 Neighbor 10.1.1.1, interface address 192.168.13.1
    In the area 0 via interface GigabitEthernet0/3/0/1
    Neighbor priority is 1, State is FULL, 10 state changes
    DR is 0.0.0.0 BDR is 0.0.0.0
    Options is 0x52
    LLS Options is 0x1 (LR)
    Dead timer due in 00:00:36
    Neighbor is up for 1w2d
    Number of DBD retrans during last exchange 0
    Index 3/3, retransmission queue length 0, number of retransmission 5
    First 0(0)/0(0) Next 0(0)/0(0)
    Last retransmission scan length is 1, maximum is 1
    Last retransmission scan time is 0 msec, maximum is 0 msec

 Neighbor 10.4.4.4, interface address 192.168.34.4
    In the area 0 via interface GigabitEthernet0/3/0/2
    Neighbor priority is 1, State is FULL, 48 state changes
    DR is 0.0.0.0 BDR is 0.0.0.0
    Options is 0x12
    LLS Options is 0x1 (LR)
    Dead timer due in 00:00:30
    Neighbor is up for 00:40:03
    Number of DBD retrans during last exchange 0
    Index 2/2, retransmission queue length 0, number of retransmission 6
    First 0(0)/0(0) Next 0(0)/0(0)
    Last retransmission scan length is 0, maximum is 1
    Last retransmission scan time is 0 msec, maximum is 0 msec
```

This table describes the significant fields shown in the display.

Table 107: show ospf neighbor detail Field Descriptions

Field	Description
Neighbor	Neighbor router ID.
interface address	IP address of the interface.
In the area	Area and interface through which the OSPF neighbor is known.

Field	Description
Neighbor priority	Router priority of neighbor and neighbor state.
State	OSPF state.
state changes	Number of state changes for this neighbor.
DR is	Neighbor ID of the designated router.
BDR is	Neighbor ID of the backup designated router.
Options	Hello packet options field contents. (E-bit only; possible values are 0 and 2; 2 indicates that the area is not a stub; 0 indicates that the area is a stub).)
LLS Options is 0x1 (LR)	Neighbor is NFS Cisco capable.
Dead timer	Time (in hh:mm:ss) to elapse before OSPF declares the neighbor dead.
Neighbor is up	Amount of time (in hh:mm:ss) that the OSPF neighbor has been up.
Number of DBD retrans	Number of re-sent database description packets.
Index	Index and the remaining lines of this command give detailed information about flooding information received from the neighbor.

The following example displays the traffic engineering support, hostname, and graceful restart and helper capabilities of neighboring connected routers, which are enabled for Segment Routing.

```
Router#show ospf neighbor capabilities
Mon Apr 29 08:27:07.289 UTC

Capabilities advertised in RI LSA
  TE - Traffic Engineering
  GR - Graceful Restart
  GH - Graceful Restart Helper

  NA - Capabilities TLV Not Advertised

Neighbors for OSPF 1

Area 1
Neighbor ID            Address              Interface       SR-Enabled    Capabilities
192.168.199.137        209.165.200.225/27       Gi100/0/0/0     Yes           TE, GR, GH
    HostName: Device_3
192.168.48.1           209.165.200.226/27        Gi100/0/0/4     No            NA
192.168.48.200         209.165.200.227/27       Gi10/0/0/0      Yes           TE, GR, GH
    HostName: Device_0

Total neighbor count: 3
```

This table describes the significant fields shown in the display.

Table 108: show ospf neighbor capabilties Field Descriptions

Field	Description
SR-Enabled	Indicates whether a neighbor router is enabled for Segment Routing or not. The valid values are Yes and No.
Capabilities	Includes the capabilities advertised in the RI LSA of a neighbor router. In general, Traffic Engineering (TE), Graceful Restart (GR), and Graceful Restart Helper (GH) capabilities are advertised. By default, routers advertise the GH capability in the RI LSA. To enable the TE and GR capabilities on neighboring routers, you require configuration at either the area or OSPF level. When Segment Routing is not enabled for a neighbor router, it does not advertise any capabilities in the RI LSA, which is indicated by NA.
Hostname	Indicates the hostname of the neighbor router. If the hostname is not configured, the RI LSA uses the default hostname of the neighbor router.
Neighbor ID	Router ID of the neighbor router.
Address	IP address of the neighbor router.
Interface	Interface IP address of the neighbor router.

The following example shows the capabilities of neighbor 192.168.199.130 listed in the summary information:

```
Router#show ospf neighbor 192.168.199.130
Mon Nov 6 19:13:08.861 IST
* Indicates MADJ interface
# Indicates Neighbor awaiting BFD session up

Neighbors for OSPF 1

 Neighbor 192.168.199.130, interface address 172.31.80.37
    In the area 0 via interface GigabitEthernet0/2/0/4
    Neighbor priority is 1, State is FULL, 6 state changes
    DR is 0.0.0.0 BDR is 0.0.0.0
    Options is 0x52
    LLS Options is 0x1 (LR)
    Dead timer due in 00:00:32
    Neighbor is up for 00:08:50
    Number of DBD retrans during last exchange 0
    Index 3/6, retransmission queue length 0, number of retransmission o
    First 0(0) /0(0) Next 0(0) /0(0)
    Last retransmission scan length is 0, maximum is 0
    Last retransmission scan time is l msec, maximum is 8 msec
    LS Ack list: NSR-sync pending 0, high water mark o
    Adjacency SIDs:
```

```
        Label: 24010,    Dynamic, Unprotected
        Label: 24011,    Dynamic,   Protected (Has backup)
    Neighbor Interface ID: 4
    Segment-Routing Enabled
      Strict SPF capable
      Start SID: 16000 Range Size: 8000
      Flex algo Enabled
    Capabilities advertised in RI LSA
      Graceful Restart capable
      Graceful Restart Helper capable
HostName: r5

Total neighbor count: 1
```

Related Commands

Command	Description
router ospf (OSPF), on page 1190	Configures an OSPF routing process.

show ospf request-list

To display the first ten link-state requests pending that the local router is making to the specified Open Shortest Path First (OSPF) neighbor and interface, use the **show ospf request-list** command in EXEC mode.

show ospf [*process-name*] [**vrf** {*vrf-name* | **all**}] [*area-id*] **request-list** [*type interface-path-id*] [*neighbor-id*]

Syntax Description

process-name	(Optional) Name that uniquely identifies an OSPF routing process. The process name is defined by the **router ospf** command. If this argument is included, only information for the specified routing process is displayed.
vrf	(Optional) Specifies an OSPF VPN routing and forwarding (VRF) instance.
vrf-name	(Optional) Name of the OSPF VRF. The *vrf-name* argument can be specified as an arbitrary string. The strings "default" and "all" are reserved VRF names.
all	(Optional) Specifies all OSPF VRF instances.
area-id	(Optional) Area ID. If you do not specify an area, all areas are displayed.
type	Interface type. For more information, use the question mark (?) online help function.
i *nterface-path-id*	Physical interface or virtual interface. 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.
neighbor-id	(Optional) IP address of the OSPF neighbor.

Command Default

All neighbors

Command Modes

EXEC

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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

You might use this command when the databases of two neighboring routers are out of synchronization or if the adjacency does not form between them. Adjacency means that the routers synchronize their databases when they discover each other.

You can look at the list to determine if one router is trying to request a particular database update. Entries that are suspended in the list usually indicate that updates are not being delivered. One possible reason for this behavior is a maximum transmission unit (MTU) mismatch between the routers.

You might also look at this list to make sure it is not corrupted. The list should refer to database entries that actually exist.

Request list information is transient and normally the lists are empty.

Task ID

Task ID	Operations
ospf	read

Examples

The following is sample output from the **show ospf request-list** command:

```
RP/0/RSP0/CPU0:router# show ospf request-list 10.0.124.4 GigabitEthernet3/0/0/0

Request Lists for OSPF pagent

 Neighbor 10.0.124.4, interface GigabitEthernet3/0/0/0 address 10.3.1.2

 Type   LS ID            ADV RTR         Seq NO        Age   Checksum
    1   192.168.58.17     192.168.58.17   0x80000012     12   0x0036f3
    2   192.168.58.68     192.168.58.17   0x80000012     12   0x00083f
```

This table describes the significant fields shown in the display.

Table 109: show ospf request-list 10.0.124.4 GigabitEthernet3/0/0/0 Field Descriptions

Field	Description
Neighbor	Specific neighbor receiving the request list from the local router.
Interface	Specific interface over which the request list is being sent.
Address	Address of the interface over which the request list is being sent.
Type	Type of link-state advertisement (LSA).
LS ID	Link-state ID of the LSA.
ADV RTR	IP address of the advertising router.
Seq NO	Sequence number of the LSA.
Age	Age of the LSA (in seconds).
Checksum	Checksum of the LSA.

Related Commands

Command	Description
router ospf (OSPF), on page 1190	Configures an OSPF routing process.

Command	Description
show ospf retransmission-list, on page 1244	Displays the first ten link-state entries in the retransmission list that the local router sends to the specified neighbor over the specified interface.

show ospf retransmission-list

To display the first ten link-state entries in the Open Shortest Path First (OSPF) retransmission list that the local router sends to the specified neighbor over the specified interface, use the **show ospf retransmission-list** command in EXEC mode.

show ospf [*process-name*] [**vrf** {*vrf-name* | **all**}] [*area-id*] **retransmission-list** [*type interface-path-id*] [*neighbor-id*]

Syntax Description

process-name	(Optional) Name that uniquely identifies an OSPF routing process. The process name is defined by the **router ospf** command. If this argument is included, only information for the specified routing process is displayed.
vrf *vrf-name* **all**	(Optional) Specifies an OSPF VPN routing and forwarding (VRF) instance. The *vrf-name* argument can be specified as an arbitrary string. The strings "default" and "all" are reserved VRF names.
area-id	(Optional) Area ID. If you do not specify an area, all areas are displayed.
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.
neighbor-id	(Optional) IP address of the OSPF neighbor.

Command Default

All neighbors

Command Modes

EXEC

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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

You might use this command when the databases of two neighboring routers are out of synchronization or if the adjacency is not forming between them. Adjacency means that the routers synchronize their databases when they discover each other.

You can look at the list to determine if one router is trying to request a particular database update. Entries that appear to be suspended in the list usually indicate that updates are not being delivered. One possible reason for this behavior is a maximum transmission unit (MTU) mismatch between the routers.

You might also look at this list to make sure it is not corrupted. The list should refer to database entries that actually exist.

Retransmission list information is transient, and normally the lists are empty.

Task ID

Task ID	Operations
ospf	read

Examples

The following is sample output from the **show ospf retransmission-list** command:

```
RP/0/RSP0/CPU0:router# show ospf retransmission-list 10.0.124.4 GigabitEthernet3/0/0/0

 Neighbor 10.0.124.4, interface GigabitEthernet3/0/0/0 address 10.3.1.2
```

This table describes the significant fields shown in the display.

Table 110: show ospf retransmission-list 10.0.124.4 GigabitEthernet3/0/0/0 Field Descriptions

Field	Description
Neighbor	Specified neighbor receiving the retransmission list from the local router.
Interface	Specified interface over which the retransmission list is being sent.
Address	Address of the interface.

Related Commands

Command	Description
router ospf (OSPF), on page 1190	Configures an OSPF routing process.
show ospf request-list, on page 1241	Displays the first ten link-state requests pending that the local router is making to the specified neighbor and interface.

show ospf routes

To display the Open Shortest Path First (OSPF) topology table, use the **show ospf routes** command in EXEC mode.

show ospf [*process-name*] [**vrf** {*vrf-name* | **all**}] **routes** [**connected** | **external** | **local**] [*prefix mask*] [*prefix/length*] **[multicast-intact]**

Syntax Description

process-name	(Optional) Name that uniquely identifies an OSPF routing process. The process name is defined by the **router ospf** command. If this argument is included, only information for the specified routing process is displayed.
vrf *vrf-name* **all**	(Optional) Specifies an OSPF VPN routing and forwarding (VRF) instance. The *vrf-name* argument can be specified as an arbitrary string. The strings "default" and "all" are reserved VRF names.
connected	(Optional) Displays connected routes.
external	(Optional) Displays routes redistributed from other protocols.
local	(Optional) Displays the local routes redistributed from the Routing Information Base (RIB).
prefix	(Optional) IP prefix, which limits output to a specific route. If the *prefix* argument is specified, either the *length* or *mask* argument is required.
mask	(Optional) IP address mask.
/ *length*	(Optional) Prefix length, which can be indicated as a slash (/) and number. For example, /8 indicates that the first eight bits in the IP prefix are network bits. If *length* is used, the slash is required.
multicast-intact	(Optional) Displays multicast intact paths.

Command Default

All route types

Command Modes

EXEC

Command History

Release	Modification
Release 3.7.2	This command was introduced.
Release 3.9.0	Asplain format for 4-byte Autonomous system numbers notation was supported. The input parameters and output were modified to display 4-byte autonomous system numbers and extended communities in either asplain or asdot notations.

Usage Guidelines

Use the **show ospf routes** command to display the OSPF private routing table (which contains only routes calculated by OSPF). If there is something wrong with a route in the RIB, then it is useful to check the OSPF copy of the route to determine if it matches the RIB contents. If it does not match, there is a synchronization

problem between OSPF and the RIB. If the routes match and the route is incorrect, OSPF has made an error in its routing calculation.

Task ID

Task ID	Operations
ospf	read

show ospf routes command output with TI-LFA information

This is sample output from the **show ospf routes** command with **backup-path** keyword that displays backup-path information, including TI-LFA:

```
RP/0/RSP0/CPU0:routersh ospf 1 routes 2.2.2.2/32 backup-path
Fri Apr  4 02:08:04.210 PDT

Topology Table for ospf 1 with ID 1.1.1.1

Codes: O - Intra area, O IA - Inter area
     O E1 - External type 1, O E2 - External type 2
     O N1 - NSSA external type 1, O N2 - NSSA external type 2

O    2.2.2.2/32, metric 3
     10.1.0.2, from 2.2.2.2, via GigabitEthernet0/0/0/7, path-id 1
         Backup path: TI-LFA, P node: 4.4.4.4, Labels: 16004, 123
            10.0.3.2, from 2.2.2.2, via GigabitEthernet0/0/0/3, protected bitmap 0x1
            Attribues: Metric: 104, SRLG Disjoint
```

This table describes the significant fields shown in the display.

Table 111: show ospf route Field Descriptions

Field	Description
O	OSPF route.
E	External Type 1 or 2 route.
N	NSSA Type 1 or 2
2.2.2.2/32	Network and subnet mask to which the local router has a route.
metric	Cost to reach network 10.3.1.0.
10.1.0.2	Next-hop router on the path to network 10.3.1.0.
from 2.2.2.2	Router ID 172.16.10.1 is the router that advertised this route.
via GigabitEthernet0/0/0/7	Packets destined for the given prefix (10.3.1.0/24) are sent over GigabitEthernet interface 0/0/0/7.
Backup path	Indicates the topology independent loop-free alternate backup path. Here, the backup path uses the P node 4.4.4.4.

Examples

The following is sample output from the **show ospf routes** command:

```
RP/0/RSP0/CPU0:router# show ospf routes

Topology Table for ospf 1 with ID 10.3.4.2

Codes:O - Intra area, O IA - Inter area
       O E1 - External type 1, O E2 - External type 2
       O N1 - NSSA external type 1, O N2 - NSSA external type 2

O E2 10.3.1.0/24, metric 1
       10.3.4.1, from 172.16.10.1, via GigabitEthernet 0/1/0/1
O    10.3.4.0/24, metric 1562
       10.3.4.2, directly connected, via GigabitEthernet 0/1/0/1
O E2 10.1.0.0/16, metric 1
       10.3.4.1, from 172.16.10.1, via GigabitEthernet 0/1/0/1
O IA 10.10.10.0/24, metric 1572
       10.3.4.1, from 172.16.10.1, via GigabitEthernet 0/1/0/1
O E2 130.10.10.0/24, metric 20
       10.3.4.1, from 172.16.10.1, via GigabitEthernet 0/1/0/1
```

This table describes the significant fields shown in the display.

Table 112: show ospf route Field Descriptions

Field	Description
O	OSPF route.
E	External Type 1 or 2 route.
N	NSSA Type 1 or 2
10.3.1.0/24	Network and subnet mask to which the local router has a route.
metric	Cost to reach network 10.3.1.0.
10.3.4.1	Next-hop router on the path to network 10.3.1.0.
from 172.16.10.1	Router ID 172.16.10.1 is the router that advertised this route.
via GigabitEthernet 0/1/0/1	Packets destined for the given prefix (10.3.1.0/24) are sent over GigabitEthernet interface 0/1/0/1.

The following is sample output from the **show ospf routes** command with a process name of 100:

```
RP/0/RSP0/CPU0:router# show ospf 100 routes

Topology Table for ospf 100 with ID 172.23.54.14

Codes:O - Intra area, O IA - Inter area
       O E1 - External type 1, O E2 - External type 2
       O N1 - NSSA external type 1, O N2 - NSSA external type 2

O    10.1.5.0/24, metric 1562
       10.1.5.14, directly connected, via GigabitEthernet 0/3/0/3
O IA 21.0.0.0/24, metric 1572
       10.1.5.12, from 172.23.54.12, via GigabitEthernet 0/3/0/3
O    10.0.0.0/24, metric 10
```

```
        10.0.0.12, directly connected, via GigabitEthernet 0/2/0/3
```

This table describes the significant fields shown in the display.

Table 113: show ospf 100 route Field Descriptions

Field	Description
O	OSPF route.
IA	Interarea route.
10.1.5.0/24	Network and subnet mask to which the local router has a route.
metric 1562	Cost to reach network 10.1.5.0.
10.1.5.14	Next-hop router on the path to network 10.1.5.0.
from 172.23.54.12	Router ID 172.23.54.12 is the router that advertised this route.
via GigabitEthernet 0/3/0/3	Packets destined for the given prefix (10.3.1.0/24) are sent over GigabitEthernet interface 0/3/0/3.

The following is sample output from the **show ospf routes** command with a prefix of 10.0.0.0 and a length of 24:

```
RP/0/RSP0/CPU0:router# show ospf routes 10.0.0.0/24

Topology Table for ospf 100 with ID 172.23.54.14

Codes:O - Intra area, O IA - Inter area
       O E1 - External type 1, O E2 - External type 2
       O N1 - NSSA external type 1, O N2 - NSSA external type 2

O IA 10.0.0.0/24, metric 1572
       10.1.5.12, from 172.23.54.12, via GigabitEthernet 0/3/0/3
```

This table describes the significant fields shown in the display.

Table 114: show ospf route 10.0.0.0/24 Field Descriptions

Field	Description
O	Route is an OSPF route.
IA	Route to network 10.0.0.0 is an interarea route.
10.0.0.0/24	Network and subnet mask to which the local router has a route.
metric 1572	Cost to reach network 10.0.0.0.
10.1.5.12	IP address of next-hop router on the path to network 10.0.0.0.
from 172.23.54.12	Router ID 172.23.54.12 is the router that advertised this route.

Field	Description
via GigabitEthernet 0/3/0/3	Packets destined for the given prefix (10.0.0.0/24) are sent over GigabitEthernet interface 0/3/0/3.

Related Commands

Command	Description
router ospf (OSPF), on page 1190	Configures an OSPF routing process.
show route, on page 1500	Displays current routes information in the Routing Information Base (RIB).
show rib opaques, on page 1483	Displays opaque data installed in the Routing Information Base (RIB).

show ospf sham-links

To display Open Shortest Path First (OSPF) sham-link information, use the **show ospf sham-links** command in EXEC mode.

show ospf [*process-name*] [**vrf** {*vrf-name* | **all**}] **sham-links**

Syntax Description

process-name	(Optional) Name that uniquely identifies an OSPF routing process. The process name is defined by the **router ospf** command. If this argument is included, only information for the specified routing process is displayed.
vrf	(Optional) Specifies an OSPF VPN routing and forwarding (VRF) instance.
vrf-name	(Optional) Name of the OSPF VRF. The *vrf-name* argument can be specified as an arbitrary string. The strings "default" and "all" are reserved VRF names.
all	(Optional) Specifies all OSPF VRF instances.

Command Default

No default behavior or values

Command Modes

EXEC

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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **show ospf sham-links** command to display OSPF sham-link information.

Task ID

Task ID	Operations
ospf	read

Examples

The following is sample output from the **show ospf sham-links** command:

```
RP/0/RSP0/CPU0:router# show ospf 1 vrf vrf_1 sham-links

Sham Links for OSPF 1, VRF vrf_1

Sham Link OSPF_SL0 to address 10.0.0.3 is up
Area 0, source address 10.0.0.1
IfIndex = 185
Run as demand circuit
DoNotAge LSA allowed., Cost of using 1
Transmit Delay is 1 sec, State POINT_TO_POINT,
Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
```

```
Hello due in 00:00:04
Adjacency State FULL (Hello suppressed)
Number of DBD retrans during last exchange 0
Index 2/2, retransmission queue length 0, number of retransmission 0
First 0(0)/0(0) Next 0(0)/0(0)
Last retransmission scan length is 0, maximum is 0
Last retransmission scan time is 0 msec, maximum is 0 msec
Keychain-based authentication enabled
Key id used is 2
```

This table describes the significant fields shown in the display.

Table 115: show ospf sham-links Field Descriptions

Field	Description
Sham Link OSPF_SL0 to address	Address of the destination endpoint of the sham link.
IfIndex	ifindex associated with the sham link.
Run as demand circuit	Sham link is treated as a demand circuit.
DoNotAge LSA allowed	DoNotAge LSAs are allowed to be flooded over the sham link.
Cost of using	Sham-link cost.
Transmit Delay	Sham-link transmit delay.
State	Sham-link interface state.
Timer intervals configured	Various sham-link interface-related timers.
Hello due in	Time before the next Hello is sent over the sham link.
Adjacency State	State of the adjacency with the neighbor over the sham link.
Number of DBD retrans during last exchange	Number of DBD retransmissions during the last exchange over the sham link.
Index	Area flood index.
retransmission queue length	Retransmission queue length on the sham link.
number of retransmission	Number of retransmissions over the sham-link interface.
First	First flood information.
Next	Next flood information.
Last retransmission scan length is	Last retransmission scan length on the sham-link interface.
maximum is	Maximum retransmission scan length on the sham-link interface.
Last retransmission scan time is	Last retransmission scan time on the sham-link interface.
maximum is 0 msec	Maximum retransmission scan time on the sham-link interface.

Field	Description
Keychain-based authentication enabled	Keychain-based authentication is enabled.
Key id used is	Key ID used.

show ospf summary-prefix

To display Open Shortest Path First (OSPF) aggregated summary address information, use the **show ospf summary-prefix** command in EXEC mode.

show ospf [*process-name*] [**vrf** {*vrf-name* | **all**}] **summary-prefix**

Syntax Description

process-name	(Optional) Name that uniquely identifies an OSPF routing process. The process name is defined by the **router ospf** command. If this argument is included, only information for the specified routing process is displayed.
vrf *vrf-name* **all**	(Optional) Specifies an OSPF VPN routing and forwarding (VRF) instance. The *vrf-name* argument can be specified as an arbitrary string. The strings "default" and "all" are reserved VRF names.

Command Default

All summary prefixes

Command Modes

EXEC

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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **show ospf summary-prefix** command if you configured summarization of external routes with the **summary-prefix** command and you want to display configured summary addresses.

Task ID

Task ID	Operations
ospf	read

Examples

The following is sample output from the **show ospf summary-prefix** command:

```
RP/0/RSP0/CPU0:router# show ospf summary-prefix

OSPF Process 1, summary-prefix

10.1.0.0/255.255.0.0 Metric 20, Type 2, Tag 0
```

This table describes the significant fields shown in the display.

Table 116: show ospf summary-prefix Field Descriptions

Field	Description
10.1.0.0/255.255.0.0	Summary address designated for a range of addresses. The IP subnet mask used for the summary route.
Metric	Metric used to advertise the summary routes.
Type	External link-state advertisements (LSA) metric type.
Tag	Tag value that can be used as a “match” value for controlling redistribution through route maps.

Related Commands

Command	Description
router ospf (OSPF), on page 1190	Configures an OSPF routing process.
summary-prefix (OSPF), on page 1269	Creates aggregate addresses for routes being redistributed from another routing protocol into the OSPF protocol.

show ospf virtual-links

To display parameters and the current state of Open Shortest Path First (OSPF) virtual links, use the **show ospf virtual-links** command in EXEC mode.

show ospf [*process-name*] [**vrf** {*vrf-name* | **all**}] **virtual-links**

Syntax Description

process-name	(Optional) Name that uniquely identifies an OSPF routing process. The process name is defined by the **router ospf** command. If this argument is included, only information for the specified routing process is displayed.
vrf *vrf-name* **all**	(Optional) Specifies an OSPF VPN routing and forwarding (VRF) instance. The *vrf-name* argument can be specified as an arbitrary string. The strings "default" and "all" are reserved VRF names.

Command Default

All virtual links

Command Modes

EXEC

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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **show ospf virtual-links** command to display useful information for debugging OSPF routing operations.

Task ID

Task ID	Operations
ospf	read

Examples

The following is sample output from the **show ospf virtual-links** command:

```
RP/0/RSP0/CPU0:router# show ospf virtual-links

Virtual Link to router 172.31.101.2 is up
Transit area 0.0.0.1, via interface GigabitEthernet 0/3/0/0, Cost of using 10
Transmit Delay is 1 sec, State POINT_TO_POINT
Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
Hello due in 0:00:08
Adjacency State FULL
```

This table describes the significant fields shown in the display.

Table 117: show ospf virtual-links Field Descriptions

Field	Description
Virtual Link to router 172.31.101.2 is up	OSPF neighbor and whether the link to that neighbor is up or down.
Transit area 0.0.0.1	Transit area through which the virtual link is formed.
via interface GigabitEthernet 0/3/0/0	Interface through which the virtual link is formed.
Cost of usingusing 10	Cost of reaching the OSPF neighbor through the virtual link.
Transmit Delay is 1 sec	Transmit delay (in seconds) on the virtual link.
State POINT_TO_POINT	State of the OSPF neighbor.
Timer intervals	Various timer intervals (in seconds) configured for the link.
Hello due in 0:00:08	When the next hello message is expected from the neighbor (in hh:mm:ss).
Adjacency State FULL	Adjacency state between the neighbors.

Related Commands

Command	Description
router ospf (OSPF), on page 1190	Configures an OSPF routing process.

show protocols (OSPF)

To display information about the OSPFv2 processes running on the router, use the **show protocols** command in EXEC mode.

show protocols [**afi-all** | **ipv4** | **ipv6**] [**all***protocol*]

Syntax Description

afi-all	(Optional) Specifies all address families.
ipv4	(Optional) Specifies an IPv4 address family.
ipv6	(Optional) Specifies an IPv6 address family.
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** • **eigrp** • **isis** • **ospf** • **rip** For the IPv6 address family, the options are: • **bgp** • **eigrp** • **isis** • **ospfv3**

Command Default

No default behavior or value

Command Modes

EXEC

Command History

Release	Modification
Release 3.7.2	This command was introduced.
Release 3.9.0	Asplain format for 4-byte Autonomous system numbers notation was supported. The input parameters and output were modified to display 4-byte autonomous system numbers and extended communities in either asplain or asdot notations.

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
ospf	read

Task ID	Operations
rib	read

Examples

The following is an OSPF configuration and the resulting **show protocols ospf** display:

```
RP/0/RSP0/CPU0:router#show running router ospf 1


router ospf 1
 router-id Loopback0
 nsf
 redistribute connected
 redistribute isis 3
 area 0
  mpls traffic-eng
  interface Loopback0
  !
  interface Loopback1
  !
  interface Loopback2
  !
  interface GigabitEthernet 0/3/0/0
  !
  interface GigabitEthernet 0/3/0/1
  !
  interface GigabitEthernet 0/3/0/2
  !
  interface GigabitEthernet 0/3/0/3
  !
 !
 mpls traffic-eng router-id Loopback0

!

RP/0/RSP0/CPU0:router# show protocols ospf
Routing Protocol OSPF 1
  Router Id: 55.55.55.55
  Distance: 110
  Non-Stop Forwarding: Enabled
  Redistribution:
    connected
    isis 3
  Area 0
    MPLS/TE enabled
    GigabitEthernet 0/3/0/3
    GigabitEthernet 0/3/0/2
    GigabitEthernet 0/3/0/1
    GigabitEthernet 0/3/0/0
    Loopback2
    Loopback0
```

This table describes the significant fields shown in the display.

Table 118: show protocols ospf Field Descriptions

Field	Description
Router Id	ID of the router for this configuration.
Distance	Administrative distance of OSPF routes relative to routes from other protocols.
Non-Stop Forwarding	Status of nonstop forwarding.
Redistribution	Lists the protocols that are being redistributed.
Area	Information about the current area including list of interfaces and the status of Multiprotocol Label Switching traffic engineering (MPLS TE).

snmp context (OSPF)

To specify an SNMP context for an OSPF instance, use the **snmp context** command in router configuration mode or in VRF configuration mode. To remove the SNMP context, use the **no** form of this command.

snmp context *context_name*
no snmp context *context_name*

Syntax Description

context_name	Specifies name of the SNMP context for OSPF instance.

Command Default

SNMP context is not specified.

Command Modes

Router configuration

VRF configuration

Command History

Release	Modification
Release 4.1.0	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.

The snmp-server commands need to be configured to perform SNMP request for the OSPF instance. Refer *SNMP Server Commands* module in *System Management Command Reference for Cisco ASR 9000 Series Routers* for information on using the **snmp-server** commands.

Note To map an SNMP context with a protocol instance, topology or VRF entity, use the **snmp-server context mapping** command. However, the **feature** option of this command does not work with OSPF protocol.

Task ID

Task ID	Operation
ospf	read, write

This example shows how to configure an SNMP context *foo* for OSPF instance *100*:

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#router ospf 100
RP/0/RSP0/CPU0:router(config-ospf)#snmp context foo
```

This example shows how to configure **snmp-server** commands to be used with the **snmp context** command:

```
RP/0/RSP0/CPU0:router(config)#snmp-server host 10.0.0.2 traps version 2c public udp-port
1620
RP/0/RSP0/CPU0:router(config)#snmp-server community public RW
RP/0/RSP0/CPU0:router(config)#snmp-server contact foo
RP/0/RSP0/CPU0:router(config)#snmp-server community-map public context foo
```

This is a sample SNMP context configuration for OSPF instance *100*:

```
snmp-server host 10.0.0.2 traps version 2c public udp-port 1620
snmp-server community public RW
snmp-server contact foo

snmp-server community-map public context foo


router ospf 100
 router-id 2.2.2.2
 bfd fast-detect
 nsf cisco
snmp context foo
 area 0
  interface Loopback1
  !
 !
 area 1
  interface GigabitEthernet0/2/0/1
   demand-circuit enable
  !
  interface GigabitEthernet0/3/0/0
  !
  interface GigabitEthernet0/3/0/1
  !
 !
!
```

Related Commands

Command	Description
snmp trap (OSPF)	Enables SNMP trap for an OSPF instance
snmp-server host	Specifies the recipient of an SNMP notification operation.
snmp-server community	Configures the community access string to permit access to the Simple Network Management Protocol (SNMP).
snmp-server contact	Sets the Simple Network Management Protocol (SNMP) system contact.
snmp-server community-map	Associates a Simple Network Management Protocol (SNMP) community with an SNMP context.

snmp trap (OSPF)

To enable SNMP trap for an OSPF instance, use the **snmp trap** command in VRF configuration mode. To disable SNMP trap for the OSPF instance, use the **no** form of this command.

snmp trap
no snmp trap

Syntax Description

This command has no keywords or arguments.

Command Default

Disabled.

Command Modes

VRF configuration

Command History

Release	Modification
Release 4.1.0	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
ospf	read, write

This example shows how to enable SNMP trap for OSPF instance *100* under VRF *vrf-1*:

```
Router#configure
Router(config)#router ospf 100
Router(config-ospf)#vrf vrf-1
Router(config-ospf-vrf)#snmp trap
```

Related Commands

Command	Description
snmp context ospf	

snmp trap rate-limit (OSPF)

To control the number of traps that OSPF sends by configuring window size and the maximum number of traps during that window, use the **snmp trap rate-limit** command in router configuration mode. To disable configuring the window size and maximum number of traps during the window, use the **no** form of this command.

snmp trap rate-limit *window-size max-num-traps*
no snmp trap rate-limit *window-size max-num-traps*

Syntax Description

window-size	Specifies the trap rate limit sliding window size.
max-num-traps	Specifies the maximum number of traps sent in window time.

Command Default

The default window size is 10 seconds and the maximum number of traps sent in the window time is 7.

Command Modes

Router configuration

Command History

Release	Modification
Release 3.9.0	This command was introduced. This command replaces the **snmp-server trap ospf rate-limit** command.

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
ospf	read,write

Examples

The following example shows how to set the trap rate limit sliding window size to 30 and the maximum number of traps sent to 100:

```
RP/0/RSP0/CPU0:router(config)#router ospf 100
RP/0/RSP0/CPU0:router(config-ospf)#snmp trap rate-limit 30 100
```

spf prefix-priority (OSPFv2)

To prioritize OSPFv2 prefix installation into the global Routing Information Base (RIB) during Shortest Path First (SPF) run, use the **spf prefix-priority** command in router configuration mode. To return to the system default value, use the **no** form of this command.

spf prefix-priority route-policy *policy-name*
no spf prefix-priority route-policy *policy-name*

Syntax Description

route-policy *policy-name*	Specifies the route policy to apply to OSPFv2 prefix prioritization.

Note
If SPF prefix prioritization is configured, /32 prefixes are no longer preferred by default. To retain the /32 prefixes in higher-priority queues, define the route-policy accordingly.

Command Default

SPF prefix prioritization is disabled.

Command Modes

OSPF router configuration

Command History

Release	Modification
Release 3.9.0	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.

SPF prefix prioritization is disabled, by default. In disabled mode, the /32 prefixes are installed into the global RIB before other prefixes.

If SPF prefix prioritization is enabled, routes are matched against the route-policy criteria and are assigned to the appropriate priority queue based on the spf-priority set. Unmatched prefixes, including the /32 prefixes, are placed in the low-priority queue.

If all /32 prefixes are desired in the high-priority queue or medium-priority queue, configure the following single route map:

```
prefix-set ospf-medium-prefixes
   0.0.0.0/0 ge 32
end-set
```

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to configure OSPFv2 SPF prefix prioritization:

```
  RP/0/RSP0/CPU0:router# configure
  RP/0/RSP0/CPU0:router(config)# prefix-set ospf-critical-prefixes
  RP/0/RSP0/CPU0:router(config-pfx)# 66.0.0.0/16
  RP/0/RSP0/CPU0:router(config-pfx)# end-set
  RP/0/RSP0/CPU0:router(config)# route-policy ospf-spf-priority
  RP/0/RSP0/CPU0:router(config-rpl)# if destination in ospf-critical-prefixes then set
spf-priority critical
  endif
  RP/0/RSP0/CPU0:router(config-rpl)# end-policy
  RP/0/RSP0/CPU0:router(config)# router ospf 1
  RP/0/RSP0/CPU0:router(config-ospf)# router-id 66.0.0.1
  RP/0/RSP0/CPU0:router(config-ospf)# spf prefix-priority route-policy ospf-spf-priority
```

Related Commands

Command	Description
prefix-set	Enters prefix set configuration mode and defines a prefix set.
route-policy (RPL)	Defines a route policy and enters route-policy configuration mode.

stub (OSPF)

To define an area as a stub area, use the **stub** command in area configuration mode. To disable this function, use the **no** form of this command.

stub [**no-summary**]
no stub

Syntax Description

no-summary	(Optional) Prevents an Area Border Router (ABR) from sending summary link advertisements into the stub area.

Command Default

No stub area is defined.

Command Modes

Area 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

You must configure the **stub** command on all routers in the stub area.

Use the **default-cost** command on the ABR of a stub area to specify the cost of the default route advertised into the stub area by the ABR.

To further reduce the number of link-state advertisements (LSAs) sent into a stub area, you can configure the **no-summary** keyword on the ABR to prevent it from sending summary LSAs (LSA Type 3) into the stub area.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to assign a default cost of 20 to stub network 10.0.0.0:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router ospf 201
RP/0/RSP0/CPU0:router(config-ospf)# area 10.0.0.0
RP/0/RSP0/CPU0:router(config-ospf-ar)# stub
RP/0/RSP0/CPU0:router(config-ospf-ar)# default-cost 20
RP/0/RSP0/CPU0:router(config-ospf-ar)# interface GigabitEthernet 0/3/0/3
```

Related Commands

Command	Description
authentication (OSPF), on page 1049	Enables authentication for an OSPF area.
default-cost (OSPF), on page 1071	Specifies a cost for the default summary route sent into a stub area.

summary-prefix (OSPF)

To create aggregate addresses for routes being redistributed from another routing protocol into the Open Shortest Path First (OSPF) protocol, use the **summary-prefix** command in the appropriate mode. To stop summarizing redistributed routes, use the **no** form of the command.

summary-prefix *address mask* [**not-advertise** | **tag** *tag*]
no summary-prefix *address mask*

Syntax Description

address	Summary address designated for a range of addresses.
mask	IP subnet mask used for the summary route.
not-advertise	(Optional) Suppresses summary routes that match the address and mask pair from being advertised.
tag *tag*	(Optional) Tag value that can be used as a "match" value for controlling redistribution through route policies.

Command Default

When this command is not used, specific addresses are created for each route from another route source being distributed into the OSPF protocol.

Command Modes

Router configuration

VRF 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **summary-prefix** command to cause an OSPF Autonomous System Boundary Router (ASBR) to advertise one external route as an aggregate for all redistributed routes that are covered by the address. This command summarizes only routes from other routing protocols that are being redistributed into OSPF.

You can use this command multiple times to summarize multiple groups of addresses. The metric used to advertise the summary is the lowest metric of all the more specific routes. This command helps reduce the size of the routing table.

If you want to summarize routes between OSPF areas, use the **range** command.

Task ID

Task ID	Operations
ospf	read, write

Examples

In the following example, summary address 10.1.0.0 includes address 10.1.1.0, 10.1.2.0, 10.1.3.0, and so on. Only the address 10.1.0.0 is advertised in an external link-state advertisement.

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router ospf 201
RP/0/RSP0/CPU0:router(config-ospf)# summary-prefix 10.1.0.0 255.255.0.0
```

Related Commands

Command	Description
range (OSPF), on page 1178	Consolidates and summarizes routes at an area boundary.

timers lsa group-pacing

To change the interval at which Open Shortest Path First (OSPF) link-state advertisements (LSAs) are collected into a group and refreshed, checksummed, or aged, use the **timers lsa group-pacing** command in the appropriate mode. To restore the default value, use the **no** form of this command.

timers lsa group-pacing *seconds*
no timers lsa group-pacing

Syntax Description

seconds	Interval (in seconds) at which LSAs are grouped and refreshed, checksummed, or aged. Range is 10 seconds to 1800 seconds.

Command Default

seconds : 240 seconds

Command Modes

Router configuration

VRF 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

OSPF LSA group pacing is enabled by default. For typical customers, the default group pacing interval for refreshing, checksumming, and aging is appropriate and you need not configure this feature.

The duration of the LSA group pacing is inversely proportional to the number of LSAs the router is handling. For example, if you have approximately 10,000 LSAs, decreasing the pacing interval would benefit you. If you have a very small database (40 to 100 LSAs), increasing the pacing interval to 10 to 20 minutes might benefit you slightly.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to change the OSPF pacing between LSA groups to 60 seconds:

```
RP/0/RSP0/CPU0:router# configure
 RP/0/RSP0/CPU0:router(config)# router ospf 1
 RP/0/RSP0/CPU0:router(config-ospf)# timers lsa group-pacing 60
```

timers lsa min-arrival

To limit the frequency that new instances of any particular Open Shortest Path First (OSPF) link-state advertisements (LSAs) can be accepted during flooding, use the **timers lsa min-arrival** command in the appropriate mode. To restore the default value, use the **no** form of this command.

timers lsa min-arrival *milliseconds*

no timers lsa min-arrival

Syntax Description

milliseconds	Minimum interval (in milliseconds) between accepting same LSA. Range is 0 to 600000 milliseconds.

Command Default

milliseconds : 100 milliseconds

Command Modes

Router configuration

VRF 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 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
ospf	read, write

Examples

The following example shows how to change the minimum interval between accepting the same LSA to 2 seconds:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router ospf 1
RP/0/RSP0/CPU0:router(config-ospf)# timers lsa min-arrival 2
```

timers throttle lsa all (OSPF)

To modify the Open Shortest Path First (OSPF) link-state advertisement (LSA) throttling, use the **timers throttle lsa all** command in the appropriate mode. To revert LSA throttling to default settings, use the **no** form of this command

timers throttle lsa all *start-interval hold-interval max-interval*
no timers throttle lsa all

Syntax Description

start-interval	Delay to generate first occurance of LSA in milliseconds. Range is 0 to 600000 milliseconds.
hold-interval	Minimum delay between originating the same LSA in milliseconds. Range is 1 to 600000 milliseconds.
max-interval	Maximum delay between originating the same LSA in milliseconds. Range is 1 to 600000 milliseconds.

Command Default

start-interval : 50 milliseconds

hold-interval : 200 milliseconds

max-interval : 5000 milliseconds

Command Modes

Router configuration

VRF 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

The *lsa-start* time is the delay before flooding the first instance of an LSA. The *lsa-hold* interval is the minimum time to elapse before flooding an updated instance of an LSA. The *lsa-max-wait* time is the maximum time that can elapse before flooding an updated instance of an LSA.

For quick convergence, use smaller times for the *lsa-start* time and *lsa-hold* interval. However, in relatively large networks, this may result in a large number of LSAs being flooded in a relatively short time. A balance with the *lsa-start* time and *lsa-hold* interval can be iteratively arrived at for the size of your network. The *lsa-max-wait* time can be used to ensure that OSPF reconverges within a reasonable amount of time.

Note LSA throttling is always enabled. You can change the timer values with the **timers throttle lsa all** command or specify the **no** keyword to revert back to the default settings.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to change the start, hold, and maximum wait interval values to 500, 1000, and 90,000 milliseconds, respectively:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router ospf 1
RP/0/RSP0/CPU0:router(config-ospf)# timers throttle lsa all 500 1000 90000
```

The following example is output from the show ospf command that displays the modified LSA throttle settings:

```
RP/0/RSP0/CPU0:router# show ospf

Routing Process "ospf 1" with ID 1.1.1.1
 Supports only single TOS(TOS0) routes
 Supports opaque LSA
 It is an area border router
 Initial SPF schedule delay 5000 msecs
 Minimum hold time between two consecutive SPFs 10000 msecs
 Maximum wait time between two consecutive SPFs 10000 msecs
 Initial LSA throttle delay 500 msecs
 Minimum hold time for LSA throttle 1000 msecs
 Maximum wait time for LSA throttle 90000 msecs
 Minimum LSA interval 1000 msecs. Minimum LSA arrival 1 secs
 Maximum number of configured interfaces 255
 Number of external LSA 0. Checksum Sum 00000000
 Number of opaque AS LSA 0. Checksum Sum 00000000
 Number of DCbitless external and opaque AS LSA 0
 Number of DoNotAge external and opaque AS LSA 0
 Number of areas in this router is 2. 2 normal 0 stub 0 nssa
 External flood list length 0
 Non-Stop Forwarding enabled
    Area BACKBONE(0) (Inactive)
        Number of interfaces in this area is 2
        SPF algorithm executed 8 times
        Number of LSA 2. Checksum Sum 0x01ba83
        Number of opaque link LSA 0. Checksum Sum 00000000
        Number of DCbitless LSA 0
        Number of indication LSA 0
        Number of DoNotAge LSA 0
        Flood list length 0
    Area 1
        Number of interfaces in this area is 1
        SPF algorithm executed 9 times
        Number of LSA 2. Checksum Sum 0x0153ea
        Number of opaque link LSA 0. Checksum Sum 00000000
        Number of DCbitless LSA 0
        Number of indication LSA 0
        Number of DoNotAge LSA 0
        Flood list length 0
```

Related Commands

Command	Description
show ospf, on page 1199	Displays generic information about OSPF routing processes.

timers throttle spf (OSPF)

To modify the Open Shortest Path First (OSPF) shortest path first (SPF) throttling, use the **timers throttle spf** command in the appropriate mode. To revert SPF throttling to default settings, use the **no** form of this command.

timers throttle spf *spf-start spf-hold spf-max-wait*
no timers throttle spf

Syntax Description

spf-start	Initial SPF schedule delay (in milliseconds). Range is 1 to 600000 milliseconds.
spf-hold	Minimum hold time (in milliseconds) between two consecutive SPF calculations. Range is 1 to 600000 milliseconds.
spf-max-wait	Maximum wait time (in milliseconds) between two consecutive SPF calculations. Range is 1 to 600000 milliseconds.

Command Default

spf-start:50 milliseconds

spf-hold: 200 milliseconds

spf-max-wait: 5000 milliseconds

Command Modes

Router configuration

VRF 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

The *spf-start* time is the delay before running SPF for the first time. The *spf-hold* interval is the minimum time to elapse between subsequent SPF runs. The *spf-max-wait* time is the maximum time that can elapse before running SPF again.

Tip Setting a low *spf-start* time and *spf-hold* time causes routing to switch to the alternate path more quickly if there is a failure; however, it consumes more CPU processing time.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to change the start, hold, and maximum wait interval values to 5, 1000, and 90000 milliseconds, respectively:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router ospf 1
RP/0/RSP0/CPU0:router(config-ospf)# timers throttle spf 5 1000 90000
```

transmit-delay (OSPF)

To set the estimated time required to send a link-state update packet on the interface, use the **transmit-delay** command in the appropriate mode. To return to the default value, use the **no** form of this command.

transmit-delay *seconds*
no transmit-delay *seconds*

Syntax Description

seconds	Time (in seconds) required to send a link-state update. Range is 1 to 65535 seconds.

Command Default

seconds: 1 second

Command Modes

Router configuration

Area configuration

Interface configuration

Virtual-link configuration

VRF configuration

Multi-area configuration

Sham-link 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Link-state advertisements (LSAs) in the update packet must have their ages incremented by the amount specified in the *seconds* argument before transmission. The value assigned should take into account the transmission and propagation delays for the interface.

If the delay is not added before transmission over a link,the time in which the LSA propagates over the link is not considered. This setting has significance only on very low-speed networks not supported in Cisco IOS XR software or on networks such as satellite circuits that incur a very long (greater than one second) delay time.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to configure a transmit delay for interface GigabitEthernet 0/3/0/0:

```
RP/0/RSP0/CPU0:router(config)# router ospf 1
RP/0/RSP0/CPU0:router(config-ospf)# area 0
RP/0/RSP0/CPU0:router(config-ospf-ar)# interface GigabitEthernet 0/3/0/0
RP/0/RSP0/CPU0:router(config-ospf-ar-if)# transmit-delay 3
```

Related Commands

Command	Description
show ospf, on page 1199	Displays general information about OSPF routing processes.

virtual-link (OSPF)

To define an Open Shortest Path First (OSPF) virtual link, use the **virtual-link** command in area configuration mode. To remove a virtual link, use the **no** form of this command.

virtual-link *router-id*
no virtual-link *router-id*

Syntax Description

router-id	Router ID associated with the virtual link neighbor. The router ID appears in the **show ospf** command display. The router ID can be any 32-bit router ID value specified in four-part, dotted-decimal notation.

Command Default

No virtual links are defined.

Command Modes

Area 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

All areas in an OSPF autonomous system must be physically connected to the backbone area (area 0). In some cases in which this physical connection is not possible, you can use a virtual link to connect to the backbone through a nonbackbone area. You can also use virtual links to connect two parts of a partitioned backbone through a nonbackbone area. The area through which you configure the virtual link, known as a transit area, must have full routing information. The transit area cannot be a stub or not-so-stubby area.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to establish a virtual link with default values for all optional parameters:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router ospf 201
RP/0/RSP0/CPU0:router(config-ospf)# area 10.0.0.0
RP/0/RSP0/CPU0:router(config-ospf-ar)# virtual-link 10.3.4.5
RP/0/RSP0/CPU0:router(config-ospf-ar-vl)#
```

The following example shows how to establish a virtual link with clear text authentication called mykey:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router ospf 201
RP/0/RSP0/CPU0:router(config-ospf)# area 10.0.0.0
RP/0/RSP0/CPU0:router(config-ospf-ar)# virtual-link 10.3.4.5
RP/0/RSP0/CPU0:router(config-ospf-ar-vl)# authentication-key 0 mykey
```

Related Commands

Command	Description
authentication (OSPF), on page 1049	Enables authentication for an OSPF area.
show ospf virtual-links, on page 1256	Displays parameters and the current state of OSPF virtual links

vrf (OSPF)

To configure an Open Shortest Path First (OSPF) VPN routing and forwarding (VRF) instance, use the **vrf** command in router configuration mode. To terminate an OSPF VRF, use the **no** form of this command.

vrf *vrf-name*
no vrf *vrf-name*

Syntax Description

vrf-name	Identifier of an OSPF VRF. The *vrf-name* argument can be specified as an arbitrary string. The strings "default" and "all" are reserved VRF names.

Command Default

No OSPF VRF is defined.

Command Modes

Router 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **vrf** command to explicitly configure a VRF. Commands configured under the VRF configuration mode (such as the **interface** [OSPF] and **authentication** commands) are automatically bound to that VRF.

To modify or remove the VRF, the *vrf-id* argument format must be the same as the format used when creating the area.



Note To remove the specified VRF from the router configuration, use the **no vrf** *vrf-id* command. The **no vrf** *vrf-id* command removes the VRF and all VRF options, such as **authentication** , **default-cost** , **nssa** , **range** , **stub** , **virtual-link** , and **interface.**

To avoid possibly having the router ID change under a VRF, explicitly configure the router ID using the **router-id** command.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to configure VRF vrf1 and GigabitEthernet interface 0/2/0/0. GigabitEthernet interface 0/2/0/0 is bound to VRF vrf1 automatically.

```
RP/0/RSP0/CPU0:router# configure
```

```
RP/0/RSP0/CPU0:router(config)# router ospf 1
RP/0/RSP0/CPU0:router(config-ospf)# vrf vrf1
RP/0/RSP0/CPU0:router(config-ospf-vrf)# area area1
RP/0/RSP0/CPU0:router(config-ospf-vrf-ar)# interface GigabitEthernet 0/2/0/0
```

Related Commands

Command	Description
router-id (OSPF), on page 1188	Configures a router ID for an OSPF process.



OSPFv3 Commands

This module describes the commands used to configure and monitor the IP Version 6 (IPv6) Open Shortest Path First Version 3 (OSPFv3) routing protocol.

For detailed information about OSPFv3 concepts, configuration tasks, and examples, see the *Implementing OSPF on* Cisco ASR 9000 Series Router module in the *Routing Configuration Guide for Cisco ASR 9000 Series Routers*.

- address-family (OSPFv3), on page 1288
- area (OSPFv3), on page 1289
- authentication (OSPFv3) , on page 1290
- auto-cost (OSPFv3), on page 1292
- capability vrf-lite (OSPFv3), on page 1294
- clear ospfv3 process, on page 1296
- clear ospfv3 redistribution, on page 1298
- clear ospfv3 routes, on page 1299
- clear ospfv3 statistics, on page 1300
- cost (OSPFv3), on page 1302
- database-filter all out (OSPFv3), on page 1304
- dead-interval (OSPFv3), on page 1306
- default-cost (OSPFv3), on page 1308
- default-information originate (OSPFv3), on page 1310
- default-metric (OSPFv3), on page 1312
- demand-circuit (OSPFv3), on page 1314
- distance ospfv3, on page 1316
- distribute-list prefix-list in, on page 1318
- distribute-list prefix-list out, on page 1320
- domain-id (OSPFv3), on page 1322
- encryption, on page 1324
- flood-reduction (OSPFv3), on page 1326
- graceful-restart (OSPFv3), on page 1328
- hello-interval (OSPFv3), on page 1329
- instance (OSPFv3), on page 1331
- interface (OSPFv3), on page 1333
- log adjacency changes (OSPFv3), on page 1335
- maximum interfaces (OSPFv3), on page 1336

- maximum paths (OSPFv3), on page 1337
- maximum redistributed-prefixes (OSPFv3), on page 1338
- mtu-ignore (OSPFv3), on page 1340
- neighbor (OSPFv3), on page 1342
- network (OSPFv3), on page 1344
- nssa (OSPFv3), on page 1346
- nsr (OSPFv3), on page 1348
- ospfv3 name-lookup, on page 1350
- packet-size (OSPFv3), on page 1351
- passive (OSPFv3), on page 1352
- priority (OSPFv3), on page 1354
- protocol shutdown (OSPFv3), on page 1356
- range (OSPFv3), on page 1357
- redistribute (OSPFv3), on page 1359
- retransmit-interval (OSPFv3), on page 1363
- router-id (OSPFv3), on page 1365
- router ospfv3, on page 1367
- sham-link (OSPFv3), on page 1368
- show ospfv3, on page 1370
- show ospfv3 border-routers, on page 1373
- show ospfv3 database, on page 1375
- show ospfv3 flood-list, on page 1387
- show ospfv3 interface, on page 1389
- show ospfv3 message-queue, on page 1392
- show ospfv3 neighbor, on page 1394
- show ospfv3 request-list, on page 1400
- show ospfv3 retransmission-list, on page 1402
- show ospfv3 routes, on page 1404
- show ospfv3 statistics rib-thread, on page 1406
- show ospfv3 summary-prefix, on page 1408
- show ospfv3 virtual-links, on page 1410
- show protocols (OSPFv3), on page 1412
- snmp context (OSPFv3), on page 1414
- snmp trap (OSPFv3), on page 1416
- snmp trap rate-limit (OSPFv3), on page 1417
- spf prefix-priority (OSPFv3), on page 1418
- stub (OSPFv3), on page 1420
- stub-router, on page 1422
- summary-prefix (OSPFv3), on page 1424
- timers lsa arrival, on page 1426
- timers pacing flood, on page 1427
- timers pacing lsa-group, on page 1429
- timers pacing retransmission , on page 1431
- timers throttle lsa all (OSPFv3), on page 1433
- timers throttle spf (OSPFv3), on page 1435
- trace (OSPFv3), on page 1437

- transmit-delay (OSPFv3), on page 1439
- virtual-link (OSPFv3), on page 1441
- vrf (OSPFv3), on page 1443

address-family (OSPFv3)

To enter address family configuration mode for Open Shortest Path First Version 3 (OSPFv3), use the **address-family** command in the router ospv3 configuration mode. To disable address family configuration mode, use the **no** form of this command.

address-family ipv6 [**unicast**]
no address-family ipv6 [**unicast**]

Syntax Description

ipv6	Specifies IP Version 6 (IPv6) address prefixes.
unicast	(Optional) Specifies unicast address prefixes.

Command Default

An address family is not specified.

Command Modes

Router ospfv3 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 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
ospf	read, write

Examples

The following example shows how to configure the OSPFv3 router process with IPv6 unicast address prefixes:

```
RP/0/RSP0/CPU0:router(config)# router ospfv3 1
RP/0/RSP0/CPU0:router(config-ospfv3)# address-family ipv6 unicast
```

area (OSPFv3)

To configure an Open Shortest Path First Version 3 (OSPFv3) area, use the **area** command in an appropriate configuration mode. To remove an OSPFv3 area, use the **no** form of this command.

area *area-id*
no area *area-id*

Syntax Description

area-id	Identifier of an OSPFv3 area. The *area-id* argument can be specified as either a decimal value or as an IPv4 address.

Command Default

No OSPFv3 areas are defined.

Command Modes

Router OSPFv3 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

An area must be explicitly configured with the **area** command.

Use the **area** command to place the router in area configuration mode (prompt: config-router-ar), from which you can configure area-specific settings. Commands configured under this mode (such as the **interface** command) are automatically bound to that area.

Note To remove the specified OSPFv3 area from the router ospfv3 configuration, use the **no area** *area-id* command. The **no area** *area-id* command removes the OSPFv3 area including all OSPFv3 area options, and all the OSPFv3 interfaces and interface options that are configured under the area.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to configure area 0 for OSPFv3 process 1. The GigabitEthernet 0/1/0/1 interface also is configured:

```
RP/0/RSP0/CPU0:router(config)# router ospfv3 1
RP/0/RSP0/CPU0:router(config-ospfv3)# area 0
RP/0/RSP0/CPU0:router(config-ospfv3-ar)# interface GigabitEthernet 0/1/0/1
```

authentication (OSPFv3)

To enable plain text, Message Digest 5 (MD5) authentication, or null authentication for an Open Shortest Path First Version 3 (OSPFv3) interface, use the **authentication** command in an appropriate configuration mode. To remove such authentication, use the **no** form of this command.

authentication {**ipsec spi** *spi-value* {**md5** | **sha1**} [**clear** | **password**] *password* | **disable**}
no authentication

Syntax Description

ipsec	Specifies IP Security (IPSec). IPSec supported only for OSPFv3.
spi *spi-value*	Specifies a security policy index (SPI) value. Range is 256 to 4294967295.
md5	Enables Message Digest 5 (MD5) authentication.
sha1	Enables SHA1 authentication.
clear	(Optional) Specifies that the key be unencrypted.
password	(Optional) Specifies that the key be encrypted using a two-way algorithm.
password	Any contiguous string that can be entered from the keyboard.
disable	Disables authentication for OSPFv3 packets.

Command Default

If this command is not specified in interface configuration mode, then the interface adopts the authentication parameter specified by the area.

If this command is not specified in area configuration mode, then the interface adopts the authentication parameter specified for the process.

If this command is not specified at any level, then the interface does not use authentication.

Command Modes

Interface configuration

Area configuration

Router configuration

Virtual-link 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **authentication** command to specify an authentication type for the interface, which overrides the authentication specified for the area to which this interface belongs. If this command is not included in the

configuration file, the authentication configured in the area to which the interface belongs is assumed (as specified by the area **authentication** command).

The authentication type and password must be the same for all OSPFv3 interfaces that are to communicate with each other through OSPFv3.

Note

- IPSec is supported only for Open Shortest Path First version 3 (OSPFv3). IPSec is supported only for Open Shortest Path First version 3 (OSPFv3).
- If OSPFv3 is configured along with IPsec authentication, then it is likely that adjacencies may flap on a Route Processor Fail Over (RPFO) even when NSR and/or Graceful Restart is enabled.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to enable MD5 authentication:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router ospfv3 201
RP/0/RSP0/CPU0:router(config-ospfv3)# router-id 10.1.1.1
RP/0/RSP0/CPU0:router(config-ospfv3)# authentication ipsec spi 500 md5
1234567890abcdef1234567890abcdef
```

auto-cost (OSPFv3)

To control how the Open Shortest Path First Version 3 (OSPFv3) protocol calculates default metrics for an interface, use the **auto-cost** command in an appropriate configuration mode. To set link cost based only on the interface type, use the **disable** form of this command. To re-enable OSPFv3 metric calculation for an interface according to the bandwidth of the interface, use the **no** form of this command.

auto-cost [**reference-bandwidth** *mbps* | **disable**]
no auto-cost [**reference-bandwidth** *mbps* | **disable**]

Syntax Description

reference-bandwidth	*mbps*	(Optional) Sets the rate in Mbps (bandwidth). Range is 1 to 4294967.
disable		(Optional) Sets the link cost based only on the interface type.

Command Default

mbps : 100 Mbps

Command Modes

Router ospfv3 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

By default OSPFv3 calculates the OSPFv3 metric for an interface according to the bandwidth of the interface.

The **no auto-cost disable** form of this command reenables OSPFv3 metric calculation for an interface according to the bandwidth of the interface.

To set link cost based only on the interface type, use the **disable** keyword.

If you have multiple links with high bandwidth, you might want to use a larger number to differentiate the cost on those links.

Recommended usage of cost configuration for all OSPFv3 configured interfaces is to be consistent: Either explicitly configure link costs (by using the **cost** command) or choose an appropriate default (by using the **auto-cost** command).

The value set by the **cost** command overrides the cost resulting from the **auto-cost** command.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to set the reference value for the auto cost to 64:

```
RP/0/RSP0/CPU0:router(config)# router ospfv3 1
RP/0/RSP0/CPU0:router(config-ospfv3)# auto-cost reference-bandwidth 64
```

Related Commands

Command	Description
cost (OSPFv3), on page 1302	Explicitly specifies the cost of the interface (network) for OSPF path calculation.

capability vrf-lite (OSPFv3)

To ignore DN bit in LSAs received from peers in the given VRF and to disable automatic ABR status in that VRF, use the **capability vrf-lite** command in OSPFv3 VRF configuration mode. To disable ignoring the DN bit in LSAs and to re-enable automatic ABR status in the VRF, use the **no** form of this command.

capability vrf-lite
no capability vrf-lite

This command has no keywords or arguments.

Command Default

Disabled.

Command Modes

OSPFv3 VRF configuration

Command History

Release	Modification
Release 4.1.0	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.

Use the **capability vrf-lite** command when routers (sometimes called multi-vrf CE routers) are directly connected through interfaces associated with a VRF, but without being connected to other PEs through the MPLS/VPN BGP Backbone.

When OSPFv3 is enabled in a VRF, the router is always an ABR. With the capability vrf-lite command enabled, the router becomes an ABR only if it is connected to area 0 (backbone area), and there are other (non-backbone) areas enabled on this router in the given VRF.

Note Routes may be re-introduced to the VPN backbone when this command is used.

Task ID

Task ID	Operation
ospf	read, write

This example shows how to enable VRF-lite capability for OSPFv3 instance *1* under VRF *vrf1*:

```
Router#configure
Router(config)#router ospfv3 1
Router(config-ospfv3)#vrf vrf1
```

```
Router(config-ospfv3-vrf)#capability vrf-lite
```

Related Commands

Command	Description
vrf (OSPFv3)	

clear ospfv3 process

To reset an Open Shortest Path First Version 3 (OSPFv3) router process without removing and reconfiguring it, use the **clear ospfv3 process** command in EXEC configuration mode.

clear ospfv3 [*process-name*] **process**

Syntax Description

process-name	(Optional) Name that uniquely identifies an OSPFv3 routing process. The process name is defined by the **router ospfv3** command. If this argument is included, only the specified routing process is affected. Otherwise, all OSPFv3 processes are reset.

Command Default

No default behavior or value

Command Modes

EXEC 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

When the OSPFv3 router process is reset, OSPFv3 releases all resources allocated, cleans up the internal database, uninstalls routes, and resets all OSPFv3 adjacencies.

Note The **clear ospfv3 process** command might clear the router ID configuration if the OSPF router ID is not explicitly configured through the router-id (OSPFv3), on page 1365 command.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to reset all OSPFv3 processes:

```
RP/0/RSP0/CPU0:router# clear ospfv3 process
```

The following example shows how to reset the OSPFv3 process 1:

```
RP/0/RSP0/CPU0:router# clear ospfv3 1 process
```

Related Commands

Command	Description
router-id (OSPFv3), on page 1365	Configures a router ID for the OSPFv3 process.

clear ospfv3 redistribution

To flush all the Type 5 and Type 7 link-state advertisements (LSAs) originated by an Open Shortest Path First Version 3 (OSPFv3) process, use the **clear ospfv3 redistribution** command in EXEC configuration mode.

clear ospfv3 [*process-name*] **redistribution**

Syntax Description

process-name	(Optional) Name that uniquely identifies an OSPFv3 routing process. The process name is defined by the **router ospfv3** command. If this argument is included, only the specified routing process is affected. Otherwise, all OSPFv3 processes are reset.

Command Default

No default behavior or value

Command Modes

EXEC 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **clear ospfv3 redistribution** command to cause the routing table to be read again. OSPFv3 regenerates and sends Type 5 and Type 7 link-state advertisements (LSAs) to its neighbors. If an unexpected route has appeared in the OSPFv3 redistribution, using this command corrects the issue.



Note Use of this command can cause a significant number of LSAs to flood the network. We recommend that you use this command with caution.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to clear all OSPFv3 redistributed routes from other protocols:

```
RP/0/RSP0/CPU0:router# clear ospfv3 redistribution
```

clear ospfv3 routes

To clear the Open Shortest Path First Version 3 (OSPFv3) internal route table, use the **clear ospfv3 routes** command in EXEC configuration mode.

clear ospfv3 [*process-name*] **routes**

Syntax Description

process-name	(Optional) Name that uniquely identifies an OSPFv3 routing process. The process name is defined by the **router ospfv3** command. If this argument is included, only the specified routing process is affected. Otherwise, all OSPFv3 processes are reset.

Command Default

No default behavior or value

Command Modes

EXEC 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **clear ospfv3 routes** command to force the internal route table to be repopulated by causing recalculation of the shortest path first (SPF) routing table. When the OSPFv3 routing table is cleared, OSPFv3 routes in the global routing table are also recalculated.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to clear all OSPFv3 routes from the OSPFv3 routing table and recomputes valid routes:

```
RP/0/RSP0/CPU0:router# clear ospfv3 routes
```

clear ospfv3 statistics

To clear the Open Shortest Path First Version 3 (OSPFv3) statistical counters, use the **clear ospfv3 statistics** command in EXEC configuration mode.

clear ospfv3 [*process-name*] **statistics** [**neighbor** [*type interface-path-id*] [*router-id*]]

Syntax Description

process-name	(Optional) Name that uniquely identifies an OSPFv3 routing process. The process name is defined by the **router ospfv3** command. If this argument is included, only the specified routing process is affected.
neighbor	(Optional) Clears counters for the specified neighbor only.
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.
router-id	(Optional) Specified router ID. This argument must be in 32-bit dotted-decimal notation, similar to an IPv4 address. This argument clears the counters of the specified neighbor only.

Command Default

No default behavior or value

Command Modes

EXEC 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **clear ospfv3 statistics** command to reset statistics so that subsequent changes are easily observed.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to clear the OSPFv3 statistical counters of all neighbors on GigabitEthernet interface 0/2/0/0:

```
RP/0/RSP0/CPU0:router# clear ospfv3 statistics neighbor GigabitEthernet 0/2/0/0
```

cost (OSPFv3)

To explicitly specify the cost of the interface (network) for OSPF path calculations, use the **cost** command in an appropriate configuration mode. To remove the cost, use the **no** form of this command.

cost *cost*
no cost

Syntax Description

cost	Unsigned integer value expressed as the link-state metric. Range is 1 to 65535.

Command Default

If this command is not specified in interface configuration mode, then the interface adopts the cost parameter specified by the area.

If this command is not specified in area configuration mode, then the interface adopts the cost parameter specified for the process.

If this command is not specified at any level, then the cost is based on the interface bandwidth, as specified by the **auto-cost** command.

Command Modes

Interface configuration

Area configuration

Router OSPFv3 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

The link-state metric is advertised as the link cost in the router link advertisement.

In general, the path cost is calculated using the following formula:

10^8 / bandwidth

If this value does not suit your network, you can use your own method of calculating path costs.

The value set by the **cost** command overrides the cost resulting from the **auto-cost** command.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to set the cost value to 65 for GigabitEthernet interface 0/1/0/1:

```
RP/0/RSP0/CPU0:router(config)# router ospfv3 201
RP/0/RSP0/CPU0:router(config-ospfv3)# area 0
RP/0/RSP0/CPU0:router(config-ospfv3-ar)# interface GigabitEthernet 0/1/0/1
RP/0/RSP0/CPU0:router(config-ospfv3-ar-if)# cost 65
```

Related Commands

Command	Description
auto-cost (OSPFv3), on page 1292	Controls how the OSPFv3 protocol calculates default metrics for an interface.

database-filter all out (OSPFv3)

To filter outgoing link-state advertisements (LSAs) to an Open Shortest Path First Version 3 (OSPFv3) interface, use the **database-filter all out** command in an appropriate configuration mode. To restore the forwarding of LSAs to the interface, use the **no** form of this command.

database-filter all out
no database-filter all out

Syntax Description

This command has no keywords or arguments.

Command Default

If this command is not specified in interface configuration mode, then the interface adopts the database filter parameter specified for the area.

If this command is not specified in area configuration mode, then the interface adopts the database filter parameter specified for the process.

If this command is not specified in router ospfv3 configuration mode, then the database filter is disabled and all outgoing LSAs are flooded to the interface.

Command Modes

Interface configuration

Area configuration

Router OSPFv3 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **database-filter all out** command to perform the same function that the **neighbor** command (with the **database-filter** keyword) performs on a neighbor basis.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to prevent flooding of OSPFv3 LSAs to neighbors reachable through GigabitEthernet interface 0/2/0/3:

```
RP/0/RSP0/CPU0:router(config)# router ospfv3 1
RP/0/RSP0/CPU0:router(config-ospfv3)# area 0
RP/0/RSP0/CPU0:router(config-ospfv3-ar)# interface GigabitEthernet 0/2/0/3
RP/0/RSP0/CPU0:router(config-ospfv3-ar-if)# database-filter all out
```

Related Commands

Command	Description
neighbor (OSPFv3), on page 1342	Configures OSPFv3 routers interconnecting to nonbroadcast networks.

dead-interval (OSPFv3)

To set the interval after which a neighbor is declared dead when no hello packets are observed, use the **dead-interval** command in an appropriate configuration mode. To return to the default time, use the **no** form of this command.

dead-interval *seconds*
no dead-interval

Syntax Description

seconds	Unsigned integer that specifies the interval (in seconds). The value must be the same for all nodes on the same network link. Range is 1 to 65535.

Command Default

If this command is not specified in interface configuration mode, then the interface adopts the dead interval parameter specified for the area.

If this command is not specified in area configuration mode, then the interface adopts the dead interval parameter specified for the process.

If this command is not specified in router ospfv3 configuration mode, then the dead interval is four times the interval set by the **hello-interval (OSPFv3)** command.

Command Modes

Interface configuration

Area configuration

Router OSPFv3 configuration

Virtual-link 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Two Open Shortest Path First Version 3 (OSPFv3) routers do not become adjacent if their dead interval values differ.

If the hello interval is configured, the dead interval value must be larger than the hello interval value. The dead interval value is usually configured four times larger than the hello interval value.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to set the OSPFv3 dead interval on GigabitEthernet interface 0/2/0/3 to 40 seconds:

```
RP/0/RSP0/CPU0:router(config)# router ospfv3 1
RP/0/RSP0/CPU0:router(config-ospfv3)# area 0
RP/0/RSP0/CPU0:router(config-ospfv3-ar)# interface GigabitEthernet 0/2/0/3
RP/0/RSP0/CPU0:router(config-ospfv3-ar-if)# dead-interval 40
```

Related Commands

Command	Description
hello-interval (OSPFv3), on page 1329	Specifies the interval between hello packets that the Cisco IOS XR software sends on the interface.

default-cost (OSPFv3)

To specify a cost for the default summary route sent into a stub area or not-so-stubby area (NSSA) for Open Shortest Path First Version 3 (OSPFv3) packets, use the **default-cost** command in area configuration mode. To remove the assigned default route cost, use the **no** form of this command.

default-cost *cost*
no default-cost

Syntax Description

cost	Cost for the default summary route used for a stub or NSSA area. The acceptable value is a 24-bit number ranging from 1 to 16777214.

Command Default

cost : 1

Command Modes

Area 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **default-cost** command only on an Area Border Router (ABR) attached to a stub or an NSSA area.

In all routers and access servers attached to the stub area, the area should be configured as a stub area using the **stub (OSPFv3)** command in the area configuration submode. Use the **default-cost** command only on an ABR attached to the stub area. The **default-cost** command provides the metric for the summary default route generated by the ABR into the stub area.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to assign a cost of 20 to the default route sent into area 10.15.0.0:

```
RP/0/RSP0/CPU0:router(config)# router ospfv3 201
RP/0/RSP0/CPU0:router(config-ospfv3)# area 10.15.0.0
RP/0/RSP0/CPU0:router(config-ospfv3-ar)# stub
RP/0/RSP0/CPU0:router(config-ospfv3-ar)# default-cost 20
RP/0/RSP0/CPU0:router(config-ospfv3-ar)# interface GigabitEthernet 0/3/0/1
```

Related Commands

Command	Description
stub (OSPFv3), on page 1420	Defines an area as a stub area.

default-information originate (OSPFv3)

To generate a default external route into an Open Shortest Path First Version 3 (OSPFv3) routing domain, use the **default-information originate** command in router ospfv3 configuration mode. To disable this feature, use the **no** form of this command.

default-information originate [**route-policy** *policy-name*] [**always**] [**metric** *metric-value*] [**metric-type** *type-value*] [**tag** *tag-value*]
no default-information originate [**route-policy** *policy-name*] [**always**] [**metric** *metric-value*] [**metric-type** *type-value*] [**tag** *tag-value*]

Syntax Description

route-policy *policy-name*	(Optional) Specifies the route policy to apply to default information origination.
always	(Optional) Always advertises the default route regardless of whether the software has a default route.
metric *metric-value*	(Optional) Specifies a metric used for generating the default route. The default metric value is 1. The value used is specific to the protocol.
metric-type *type-value*	(Optional) Specifies an external link type associated with the default route advertised into the OSPFv3 routing domain. It can be one of the following values: **1**—Type 1 external route **2**—Type 2 external route
tag *tag-value*	(Optional) 32-bit dotted-decimal value attached to each external route. This is not used by the OSPFv3 protocol itself. It may be used to communicate information between autonomous system boundary routers (ASBRs). If a tag is not specified, then zero (0) is used.

Command Default

A default external route into an OSPFv3 routing domain is not generated.

metric-value : 1

type-value : Type 2

tag-value: 0

Command Modes

Router ospfv3 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Whenever you use the **redistribute** or **default-information** command to redistribute routes into an OSPFv3 routing domain, the software automatically becomes an ASBR. However, an ASBR does not, by default,

generate a default route into the OSPFv3 routing domain. The software still must have a default route for itself before it generates one, except when you have specified the **always** keyword.

The **default-information originate** route-policy attach point conditionally injects the default route 0.0.0.0/0 into the OSPF link-state database, and is done by evaluating the attached policy. If any routes specified in the policy exist in the global RIB, then the default route is inserted into the link-state database. If there is no match condition specified in the policy, the policy passes and the default route is generated into the link-state database.

For information about the default-information originate attach point, see the *OSPF v3 Policy Attach Points* section in the *Implementing Routing Policy* chapter in *Routing Configuration Guide for Cisco ASR 9000 Series Routers*.

When you use the **default-information originate** command for the OSPFv3 process, the default network must reside in the routing table.

For information about routing policies, see the *Routing Policy Commands* chapter in the *Routing Command Reference for Cisco ASR 9000 Series Routers*.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to specify a metric of 100 for the default route redistributed into the OSPFv3 routing domain and an external metric type of Type 1:

```
RP/0/RSP0/CPU0:router(config)#router ospfv3 109
RP/0/RSP0/CPU0:router(config-ospfv3)#default-information originate metric 100 metric-type 1
```

Related Commands

Command	Description
redistribute (OSPFv3), on page 1359	Redistributes routes from one routing domain into another routing domain.

default-metric (OSPFv3)

To set default metric values for routes redistributed from another protocol into Open Shortest Path First Version 3 (OSPFv3), use the **default-metric** command in an appropriate configuration mode. To return to the default state, 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.

Command Default

Built-in, automatic metric translations, as appropriate for each routing protocol

Command Modes

Router OSPFv3 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **default-metric** command with the **redistribute** command to cause the current routing protocol to use the same metric value for all redistributed routes. A default metric helps solve the problem of redistributing routes with incompatible metrics. Whenever metrics do not convert, use a default metric to provide a reasonable substitute and enable the redistribution to proceed.

The default-metric value configured in OSPF configuration does not apply to connected routes that are redistributed to OSPF using the **redistribute connected** command. To set a non-default metric for connected routes, configure OSPF with the **redistribute connected metric** *metric-value* command.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to configure a router with both the Intermediate System-to-Intermediate System (IS-IS) and the OSPFv3 routing protocols. The OSPFv3 routing protocol advertises IS-IS derived routes and assigns the routes a metric of 10:

```
RP/0/RSP0/CPU0:router(config)# router ospfv3 1
RP/0/RSP0/CPU0:router(config-ospfv3)# default-metric 10
RP/0/RSP0/CPU0:router(config-ospfv3)# redistribute isis IS-IS_isp
```

Related Commands

Command	Description
redistribute (OSPFv3), on page 1359	Redistributes routes from one routing domain into another routing domain.

demand-circuit (OSPFv3)

To configure the Open Shortest Path First Version 3 (OSPFv3) router process to treat the interface as an OSPFv3 demand circuit, use the **demand-circuit** command in an appropriate configuration mode. To remove the demand circuit designation from the interface, use the **no** form of this command.

demand-circuit [**disable**]
no demand-circuit

Syntax Description

disable	(Optional) Disables the demand circuit configuration that may have been specified at a higher level in the configuration.

Command Default

If this command is not specified in interface configuration mode, then the interface adopts the demand circuit parameter specified for the area.

If this command is not specified in area configuration mode, then the interface adopts the demand circuit parameter specified for the process.

If this command is not specified at any level, then the interface is not a demand circuit.

Command Modes

Interface configuration

Area configuration

Router OSPFv3 configuration

Virtual-link 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

On point-to-point interfaces, only one end of the demand circuit must be configured with the **demand-circuit** command. Periodic hello messages are suppressed and periodic refreshes of link-state advertisements (LSAs) do not flood the demand circuit. This command allows the underlying data link layer to be closed when the topology is stable. In point-to-multipoint topology, only the multipoint end must be configured with this command.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to configure GigabitEthernet interface 0/3/0/1 as an on-demand circuit:

```
RP/0/RSP0/CPU0:router(config)# router ospfv3 1
RP/0/RSP0/CPU0:router(config-ospfv3)# area 0
RP/0/RSP0/CPU0:router(config-ospfv3)# interface GigabitEthernet 0/3/0/1
RP/0/RSP0/CPU0:router(config-ospfv3-if)# demand-circuit
```

distance ospfv3

To define the Open Shortest Path First Version 3 (OSPFv3) route administrative distances based on route type, use the **distance ospfv3** command in router ospfv3 configuration mode. To restore the default value, use the **no** form of this command.

distance ospfv3 {**intra-area** | **inter-area** | **external**} *distance*
no distance ospfv3

Syntax Description

intra-area | **inter-area** | **external**	Type of area. It can be one of the following values: **intra-area** —All routes within an area. **inter-area** —All routes from one area to another area. **external** —All routes from other routing domains, learned by redistribution.
distance	The route administrative distance.

Command Default

distance : 110

Command Modes

Router ospfv3 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

You must specify one of the keywords.

Use the **distance ospfv3** command to perform the same function as the **distance** command used with an access list. However, the **distance ospfv3** command sets a distance for an entire group of routes, rather than a specific route that passes an access list.

A common reason to use the **distance ospfv3** command is when you have multiple OSPFv3 processes with mutual redistribution, and you want to prefer internal routes from one over external routes from the other.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to change the external distance to 200, making it less reliable:

```
RP/0/RSP0/CPU0:router(config)# router ospfv3 1
RP/0/RSP0/CPU0:router(config-ospfv3)# redistribute ospfv3 2
```

```
RP/0/RSP0/CPU0:router(config-ospfv3)# distance ospfv3 external 200
RP/0/RSP0/CPU0:router(config-ospfv3)# exit
RP/0/RSP0/CPU0:router(config)# router ospfv3 2
RP/0/RSP0/CPU0:router(config-ospfv3)# redistribute ospfv3 1
RP/0/RSP0/CPU0:router(config-ospfv3)# distance ospfv3 external 200
```

Related Commands

Command	Description
distance ospf	Defines an administrative distance.

distribute-list prefix-list in

To filter the routes that Open Shortest Path First Version 3 (OSPFv3) installs in the Routing Information Base (RIB), use the **distribute-list prefix-list in** command in an appropriate configuration mode. To remove the filter, use the **no** form of this command.

distribute-list prefix-list *prefix-list-name* **in**
no distribute-list prefix-list *prefix-list-name* **in**

Syntax Description

prefix-list-name	IP Version 6 (IPv6) prefix list name. The list defines which IPv6 prefixes are installed in the RIB.

Command Default

All routes learned by OSPFv3 are installed in the RIB.

Command Modes

Interface configuration

Router OSPFv3 configuration

Command History

Release	Modification
Release 3.7.2	This command was introduced.
Release 3.9.0	Asplain format for 4-byte Autonomous system numbers notation was supported. The input parameters and output were modified to display 4-byte autonomous system numbers and extended communities in either asplain or asdot notations.

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.

Use the **distribute-list prefix-list** command to limit the routes that OSPFv3 installs in the RIB of your router. This command does not affect the information sent to other OSPFv3 routers or the routes that these routers compute and install.

Note Because the other OSPFv3 routers are not aware of any omissions in the RIB, they may send traffic addressed to the missing prefixes. If no other provision has been made for these prefixes, the packets are dropped.

When this command is specified in router ospfv3 configuration mode, the filter applies to all routes computed by OSPFv3.

When this command is specified in interface configuration mode, the filter applies only to routes that forward outgoing traffic over that interface.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to prevent OSPFv3 from installing any routes that have 2001:e624 as the first 32 bits of the address. OSPFv3 is also prevented from installing routes to 2002::/16 that use GigabitEthernet interface 0/2/0/0 as the next-hop interface:

```
RP/0/RSP0/CPU0:router(config)# ipv6 prefix-list preflist1
RP/0/RSP0/CPU0:router(config-ipv6-pfx)# deny 2001:e624::/32 le 128
RP/0/RSP0/CPU0:router(config-ipv6-pfx)# permit ::/0 le 128
!
RP/0/RSP0/CPU0:router(config)# ipv6 prefix-list preflist2
RP/0/RSP0/CPU0:router(config-ipv6-pfx)# deny 2002::/16
RP/0/RSP0/CPU0:router(config-ipv6-pfx)# permit ::/0 le 128
!
RP/0/RSP0/CPU0:router(config)# router ospfv3 1
RP/0/RSP0/CPU0:router(config-ospfv3)# distribute-list prefix-list preflist1 in
RP/0/RSP0/CPU0:router(config-ospfv3)# area 1
RP/0/RSP0/CPU0:router(config-ospfv3-ar)# interface GigabitEthernet 0/2/0/0
RP/0/RSP0/CPU0:router(config-ospfv3-ar-if)# distribute-list prefix-list preflist2 in
```

distribute-list prefix-list out

To filter the routes redistributed into Open Shortest Path First Version 3 (OSPFv3) from other routing protocols, use the **distribute-list prefix-list out** command in an appropriate configuration mode. To remove the filter, use the **no** form of this command.

distribute-list prefix-list *prefix-list-name* **out** [*protocol* [*process-id*]]
no distribute-list prefix-list *prefix-list-name* **out** [*protocol* [*process-id*]]

Syntax Description

prefix-list-name	IP Version 6 (IPv6) prefix list name. The list defines which IPv6 prefixes are installed in the RIB.
protocol	(Optional) Source protocol from which routes are being redistributed. It can be one of the following keywords: **bgp, eigrp, isis**, **ospfv3, static**, and **connected**. The **static** keyword is used to redistribute IPv6 static routes. The **connected** keyword refers to routes that are established automatically because IPv6 is enabled on an interface. For routing protocols such as OSPFv3 and Intermediate System-to-Intermediate System (IS-IS), these routes are redistributed as external to the autonomous system.
process-id	(Optional) For the **bgp** keyword, an autonomous system number has the following ranges: • 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. For the **eigrp** keyword, an autonomous system number. For the **isis** keyword, an optional argument that defines a meaningful name for a routing process. You can specify only one IS-IS process for each router. Creating a name for a routing process means that you use names when configuring routing. For the **ospfv3** keyword, an appropriate OSPFv3 process name from which routes are to be redistributed. The value takes the form of a string. A decimal number can be entered, but it is stored internally as a string.

Command Default

All routes from protocols specified in the redistribute (OSPFv3), on page 1359 command are redistributed into OSPFv3.

Command Modes

Router OSPFv3 configuration

Command History

Release	Modification
Release 3.7.2	This command was introduced.
Release 3.9.0	Asplain format for 4-byte Autonomous system numbers notation was supported.

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.

Routes may be redistributed into OSPFv3 from several other routing protocols or from other OSPFv3 processes. These routes are then communicated to other OSPFv3 routes through Type 5 (External) or Type 7 not-so-stubby area (NSSA) link-state advertisements (LSAs). Use the **distribute-list prefix-list out** command to control redistribution by matching redistributed routes against an IPv6 prefix list. Only routes permitted by the prefix list are redistributed into OSPFv3.

Each protocol being redistributed into OSPFv3 can have a separate prefix list. In addition, a prefix list can be defined that applies to all protocols.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to prevent OSPFv3 from redistributing routes that have 2001:e624 as the first 32 bits of the address. In addition, routes with a prefix beginning with 2064 are not redistributed from Border Gateway Protocol (BGP) autonomous system 1, and only those routes are redistributed from BGP autonomous system 5.

```
RP/0/RSP0/CPU0:router(config)# ipv6 prefix-list p1
RP/0/RSP0/CPU0:router(config-ipv6-pfx)# deny 2001:e624::/32 le 128
RP/0/RSP0/CPU0:router(config-ipv6-pfx)# permit ::/0 le 128
!
RP/0/RSP0/CPU0:router(config)# ipv6 prefix-list p2
RP/0/RSP0/CPU0:router(config-ipv6-pfx)# deny 2064::/16 le 128
RP/0/RSP0/CPU0:router(config-ipv6-pfx)# permit ::/0 le 128
!
RP/0/RSP0/CPU0:router(config)# ipv6 prefix-list p3
RP/0/RSP0/CPU0:router(config-ipv6-pfx)# permit 2064::/16 le 128
!
RP/0/RSP0/CPU0:router(config)# router ospfv3 1
RP/0/RSP0/CPU0:router(config-ospfv3)# redistribute bgp 1
RP/0/RSP0/CPU0:router(config-ospfv3)# redistribute bgp 5
RP/0/RSP0/CPU0:router(config-ospfv3)# distribute-list prefix-list p1 out
RP/0/RSP0/CPU0:router(config-ospfv3)# distribute-list prefix-list p2 out bgp 1
RP/0/RSP0/CPU0:router(config-ospfv3)# distribute-list prefix-list p3 out bgp 5
```

Related Commands

Command	Description
redistribute (OSPFv3), on page 1359	Redistributes routes from one routing domain into another routing domain for OSPFv3.

domain-id (OSPFv3)

To specify the Open Shortest Path First Version 3 (OSPFv3) VPN routing and forwarding (VRF) domain ID, use the **domain-id** command in VRF configuration mode. To remove the OSPFv3 VRF domain ID, use the **no** form of this command.

domain-id [**secondary**] **type** [**0005** | **0105** | **0205**] **value** *domain-id_value*
no domain-id [**secondary**] **type** [**0005** | **0105** | **0205**] **value** *domain-id-value*

Syntax Description

secondary	(Optional) OSPFv3 secondary domain ID.
type	Primary OSPFv3 domain ID in hexadecimal format. • 0005 —Type 0x0005 • 0105 —Type 0x0105 • 0205—Type 0x0205
value	OSPF domain ID value in hexadecimal format.
domain-id-value	OSPF domain ID extended community value as a 6 byte hexadecimal number.

Command Default

No domain ID is specified.

Command Modes

VRF configuration

Command History

Release	Modification
Release 4.1.0	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.

If no value is specified for the domain ID, the default is Null (all zeros) primary domain ID. One or more secondary domain IDs can be specified.

Task ID

Task ID	Operation
ospf	read, write

This example shows how to specify a domain ID with type *0105* and value *AABBCCDDEEFF*:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router ospf 1
```

```
RP/0/RSP0/CPU0:router(config-ospf)# vrf vrf_1
RP/0/RSP0/CPU0:router(config-ospf-vrf)# domain-id type 0105 value AABBCCDDEEFF
```

Related Commands

Command	Description
vrf (OSPFv3), on page 1443	Enters OSPFv3 configuration submode

encryption

To encrypt and authenticate Open Shortest Path First Version 3 (OSPFv3) packets, use the **encryption** command in an appropriate configuration mode. To remove the encryption, use the **no** form of this command.

encryption {**disable** | **ipsec spi** *spi-value* **esp** {**3des** | **aes** [**192** | **256**] | **des** | **null** [**clear** | **password**] *encrypt-password*} [**authentication** {**md5** | **sha1**} [**clear** | **password**] *auth-password*]}
no encryption

Syntax Description

disable	Disables OSPFv3 packet encryption.
ipsec spi	Specifies IPSec ESP encryption and authentication with the Security Parameter Index (SPI) value. IPSec supported only for OSPFv3.
spi-value	SPI value. Range is 256 to 4294967295.
esp	Specifies Encryption Service Payload (ESP) encryption parameters.
3des	Specifies the triple DES algorithm.
aes	Specifies the Advanced Encryption Standard (AES) algorithm.
192	(Optional) Specifies the 192-bit AES algorithm
256	(Optional) Specifies the 256-bit AES algorithm
des	Specifies the Data Encryption Standard (DES) algorithm.
null	Specifies no AES algorithm.
md5	Enables Message Digest 5 (MD5) authentication.
sha1	Enables SHA1 authentication.
clear	Specifies that the key be unencrypted.
password	Specifies that the key be encrypted using a two-way algorithm.
encrypt-password	Any contiguous string that can be entered from the keyboard as the encryption password.
auth-password	Any contiguous string that can be entered from the keyboard as the authentication password.

Command Default

No default behavior or values.

Command Modes

Interface configuration

Router OSPFv3 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **encryption** command to encrypt and authenticate OSPFv3 packets.

Note

- IPSec is supported only for Open Shortest Path First version 3 (OSPFv3).
- If OSPFv3 is configured along with IPsec authentication, then it is likely that adjacencies may flap on a Route Processor Fail Over (RPFO) even when NSR and/or Graceful Restart is enabled.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to encrypt and authenticate OSPFv3 packets:

```
RP/0/RSP0/CPU0:router(config)#router ospfv3 1
RP/0/RSP0/CPU0:router(config-ospfv3)#encryption ipsec spi 256 esp 3des clear
```

flood-reduction (OSPFv3)

To suppress the unnecessary flooding of link-state advertisements (LSAs) in stable topologies, use the **flood-reduction** command in an appropriate configuration mode. To disable this feature, use the **no** form of this command.

flood-reduction [**disable**]
no flood-reduction

Syntax Description

disable	(Optional) Turns off this functionality at a specific level.

Note
The **disable** keyword is not available in router ospfv3 configuration mode.

Command Default

If this command is not specified in interface configuration mode, then the interface adopts the flood reduction parameter specified by area.

If this command is not specified in area configuration mode, then the interface adopts the flood reduction parameter specified for the process.

If this command is not specified at any level, then flood reduction is disabled.

Command Modes

Interface configuration

Area configuration

Router OSPFv3 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

All routers that support Open Shortest Path First Version 3 (OSPFv3) demand circuits are compatible with routers supporting flooding reduction.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to reduce the flooding of unnecessary LSAs for area 0:

```
RP/0/RSP0/CPU0:router(config)# router ospfv3 1
RP/0/RSP0/CPU0:router(config-ospfv3)# area 0
```

```
RP/0/RSP0/CPU0:router(config-ospfv3-ar)# interface GigabitEthernet 0/1/0/3
RP/0/RSP0/CPU0:router(config-ospfv3-ar-if)# flood-reduction
```

Related Commands

Command	Description
show ospfv3 interface, on page 1389	Displays OSPFv3-related interface information.
show ospfv3 neighbor, on page 1394	Displays OSPFv3 neighbor information on an individual interface basis.

graceful-restart (OSPFv3)

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

graceful-restart [**helper disable** | **interval** *interval* | **lifetime** *lifetime*]
no graceful-restart [**helper disable** | **interval** *interval* | **lifetime** *lifetime*]

Syntax Description

helper disable	(Optional) Disables the routers helper support level.
interval *interval*	(Optional) Specifies the minimum interval between graceful restarts. Range is 90 to 3600 seconds.
lifetime *lifetime*	(Optional) Specifies the maximum route lifetime following a restart. Range is 90 to 3600 seconds.

Command Default

No default behavior or values.

Command Modes

Router OSPFv3 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 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
ospf	read, write

Examples

The following example shows how to enable the Graceful Restart feature with a minimum interval between restarts of 300 seconds:

```
RP/0/RSP0/CPU0:router(config)# router ospfv3 1
RP/0/RSP0/CPU0:router(config-ospfv3)# graceful-restart interval 300
```

Related Commands

Command	Description
show ospfv3 interface, on page 1389	Displays OSPFv3-related interface information.
show ospfv3 neighbor, on page 1394	Displays OSPFv3 neighbor information on an individual interface basis.

hello-interval (OSPFv3)

To specify the interval between hello packets that Open Shortest Path First Version 3 (OSPFv3) sends on an interface, use the **hello-interval** command in an appropriate configuration mode. To return to the default time, use the **no** form of this command.

hello-interval *seconds*
no hello-interval

Syntax Description

seconds	Interval (in seconds). The value must be the same for all nodes on a specific network.

Command Default

If this command is not specified in interface configuration mode, then the interface adopts the hello interval parameter specified by area.

If this command is not specified in area configuration mode, then the interface adopts the hello interval parameter specified for the process.

If this command is not specified at any level, then the hello interval is 10 seconds (broadcast) or 30 seconds (non-broadcast).

Command Modes

Interface configuration

Area configuration

Router OSPFv3 configuration

Virtual-link 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

The hello interval value is advertised in the hello packets. The shorter the hello interval, the faster topological changes are detected, but more routing traffic ensues. This value must be the same for all routers and access servers on a specific network.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to set the interval between hello packets to 15 seconds on GigabitEthernet interface 0/3/0/2:

```
RP/0/RSP0/CPU0:router(config)# router ospfv3 1
RP/0/RSP0/CPU0:router(config-ospfv3)# area 0
RP/0/RSP0/CPU0:router(config-ospfv3-ar)# interface GigabitEthernet 0/3/0/2
RP/0/RSP0/CPU0:router(config-ospfv3-ar-if)# hello-interval 15
```

Related Commands

Command	Description
dead-interval (OSPFv3), on page 1306	Sets the interval after which a neighbor is declared dead when no hello packets are observed.

instance (OSPFv3)

To set the 8-bit instance ID used in Open Shortest Path First Version 3 (OSPFv3) packets sent on an interface, use the **instance** command in an appropriate configuration mode. To remove the instance ID, use the **no** form of this command.

instance *instance-id*
no instance *instance-id*

Syntax Description

instance-id	Instance identifier sent in OSPFv3 packets. Range is 0 to 255. The same value must be used by all the communicating OSPFv3 routers on a link.

Command Default

If this command is not specified in interface configuration mode, then the interface adopts the instance parameter specified by the area.

If this command is not specified in area configuration mode, then the interface adopts the instance parameter specified for the process.

If this command is not specified at any level, then the instance is 0.

Command Modes

Interface configuration

Area configuration

Router OSPFv3 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

The OSPFv3 routing protocol allows multiple unrelated OSPFv3 processes to share a link by using an 8-bit "instance" value to demultiplex the protocol packets. Each OSPFv3 process sets its configured instance value in the OSPFv3 packets that it sends, and ignores received packets with instance values from other OSPFv3 processes.



Note The *instance-id* argument should not be confused with the *process-name* argument that is specified by the **router ospfv3** command. The former is an 8-bit integer that is sent to other routers as part of the OSPFv3 protocol, and the latter is a 1- to 40-character ASCII string that is significant only within a given router. The instance ID value is also unrelated to the router ID that is specified by the **router-id** command, which is a 32-bit integer value that uniquely identifies a router within an OSPFv3 routing domain.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to set the instance value for GigabitEthernet interface 0/3/0/1 to 42:

```
RP/0/RSP0/CPU0:router(config)# router ospfv3 1
RP/0/RSP0/CPU0:router(config-ospfv3)# area 0
RP/0/RSP0/CPU0:router(config-ospfv3-ar)# interface GigabitEthernet 0/3/0/1
RP/0/RSP0/CPU0:router(config-ospfv3-ar-if)# instance 42
```

Related Commands

Command	Description
router ospfv3, on page 1367	Configures an OSPFv3 routing process.
router-id (OSPFv3), on page 1365	Configures a router ID for the OSPFv3 routing process.

interface (OSPFv3)

To define the interfaces on which Open Shortest Path First Version 3 (OSPFv3) runs, use the **interface** command in an appropriate configuration mode. To disable OSPFv3 routing for an interface, use the **no** form of this command.

interface *type interface-path-id*
no interface *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

An interface is not defined.

Command Modes

Area 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **interface** command to associate a specific interface with an OSPFv3 area. The interface remains associated with the area even when the IPv6 address of the interface changes.

Similar to IPv4 address behavior for the **interface** command, all configured IPv6 addresses are advertised on an interface after the interface is associated to the OSPF routing process. The only difference is, IPv6 addresses can have multiple primary addresses.

This command places the router in interface configuration mode (prompt: config-router-ar-if), from which you can configure interface-specific settings. Commands configured under this mode (such as the **cost** command) are automatically bound to that interface.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to define two interfaces that belong to area 1. The cost value for packets on GigabitEthernet interface 0/3/0/1 is set at 40; the cost value for GigabitEthernet interface 0/3/0/2 is 65:

```
RP/0/RSP0/CPU0:router(config)# router ospfv3 1
RP/0/RSP0/CPU0:router(config-ospfv3)# area 1
RP/0/RSP0/CPU0:router(config-ospfv3-ar)# interface GigabitEthernet 0/3/0/1
RP/0/RSP0/CPU0:router(config-ospfv3-ar-if)# cost 40
RP/0/RSP0/CPU0:router(config-ospfv3-ar-if)# exit
RP/0/RSP0/CPU0:router(config-ospfv3-ar)# interface GigabitEthernet 0/3/0/2
RP/0/RSP0/CPU0:router(config-ospfv3-ar-if)# cost 65
RP/0/RSP0/CPU0:router(config-ospfv3-ar-if)# exit
```

log adjacency changes (OSPFv3)

To change the default syslog messages for Open Shortest Path First Version 3 (OSPFv3) neighbor state changes, use the **log adjacency changes** command in an appropriate configuration mode. To suppress all adjacency change messages, use the **disable** keyword.

log adjacency changes [**detail** | **disable**]

Syntax Description

detail	(Optional) Provides all (DOWN, INIT, 2WAY, EXSTART, EXCHANGE, LOADING, FULL) adjacency state changes.
disable	(Optional) Disables the neighbor state change messages.

Command Default

Neighbor state change messages are enabled.

Command Modes

Router OSPFv3 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

By default, you are notified of OSPFv3 neighbor changes without explicitly configuring the **log adjacency changes** command. The syslog messages that are sent provide a high-level view of changes to the state of the OSPFv3 peer relationship.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to disable neighbor state change messages:

```
RP/0/RSP0/CPU0:router(config)# router ospfv3 1
RP/0/RSP0/CPU0:router(config-ospfv3)# log adjacency changes disable
```

The following example shows how to re-enable syslog messages for any OSPFv3 neighbor state changes:

```
RP/0/RSP0/CPU0:router(config)# router ospfv3 1
RP/0/RSP0/CPU0:router(config-ospfv3)# log adjacency changes
```

maximum interfaces (OSPFv3)

To control the maximum number of interfaces that can be configured under an Open Shortest Path First Version 3 (OSPFv3) process, use the **maximum interfaces** command in an appropriate configuration mode. To remove the **maximum interfaces** command from the configuration file and restore the system to its default condition with respect to the routing protocol, use the **no** form of this command.

maximum interfaces *number-interfaces*
no maximum interfaces

Syntax Description

number-interfaces	Maximum number of interfaces that can be configured for this OSPFv3 process. Range is 1 to 4294967295.

Command Default

If the command is not specified, the default is 1024.

Command Modes

Router OSPFv3 configuration

Command History

Release	Modification
Release 3.7.2	This command was introduced.
Release 4.1.1	The range for number of interfaces was changed to 1 to 4294967295 from 1 to 1024. The default number of interfaces was changed to 1024 from 255.

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
ospf	read, write

Examples

This example shows how to allow a maximum of 1500 interfaces in an OSPFv3 process:

```
RP/0/RSP0/CPU0:router(config)# router ospfv3 1
RP/0/RSP0/CPU0:router(config-ospfv3)# maximum interfaces 1500
```

maximum paths (OSPFv3)

To control the maximum number of parallel routes that the Open Shortest Path First Version 3 (OSPFv3) can support, use the **maximum paths** command in an appropriate configuration mode. To remove the **maximum paths** command from the configuration file and restore the system to its default condition with respect to the routing protocol, use the **no** form of this command.

maximum paths *maximum-routes-number*
no maximum paths

Syntax Description

maximum-routes-number	Maximum number of parallel routes that OSPFv3 can install in a routing table. Range is 1 to 32. **Note** The maximum number of paths that can be configured is 32.

Command Default

32 paths

Command Modes

Router OSPFv3 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

When the maximum number of parallel routes is reduced, all existing paths are pruned and paths reinstalled at the new maximum number. During this route-reduction period, you may experience some packet loss for a few seconds. This may impact route traffic.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to allow a maximum of two paths to a destination:

```
RP/0/RSP0/CPU0:router(config)# router ospfv3 1
RP/0/RSP0/CPU0:router(config-ospfv3)# maximum paths 2
```

maximum redistributed-prefixes (OSPFv3)

To limit the number of prefixes redistributed into Open Shortest Path First Version 3 (OSPFv3) or to generate a warning when the number of prefixes redistributed into OSPFv3 reaches a maximum, use the **maximum redistributed-prefixes** command in an appropriate configuration mode. To remove the values, use the **no** form of this command.

maximum redistributed-prefixes *limit* [*threshold*] [**warning-only**]
no maximum redistributed-prefixes

Syntax Description

limit	Maximum number of IP Version 6 (IPv6) prefixes that are allowed to be redistributed into OSPFv3, or, if the **warning-only** keyword is present, sets the number of prefixes allowed to be redistributed into OSPFv3 before the system logs a warning message. Range is 1 to 4294967295. **Note** If the **warning-only** keyword is also configured, this value does not limit redistribution; it is simply the number of redistributed prefixes that, when reached, causes a warning message to be logged.
threshold	(Optional) Percentage of the value set for the maximum number of redistributed prefixes that, when reached, causes a warning message to be logged.
warning-only	(Optional) Causes a warning to be logged when the number of routes defined by the *limit* argument have been redistributed. Additional redistribution is not prevented.

Command Default

limit : 10240

threshold : 75 percent

Command Modes

Router OSPFv3 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

If someone mistakenly injects a large number of IPv6 routes into OSPFv3, perhaps by redistributing Border Gateway Protocol (BGP) into OSPFv3, the network can be severely flooded. Limiting the number of redistributed routes prevents this potential problem.

When the **maximum redistributed-prefixes** command is configured, if the number of redistributed routes reaches the maximum value configured, no more routes are redistributed (unless the **warning-only** keyword is configured).

The redistribution limit applies only to external IPv6 prefixes. Default routes and summarized routes are not limited.

The limit is tracked separately for each not-so-stubby-area (NSSA) because redistribution to NSSAs is done independently for each NSSA and independently of all other regular areas.

Select a maximum value based on your knowledge of how many prefixes are redistributed on the router to the OSPFv3 process.

Task ID

Task ID	Operations
ospf	read, write

Examples

This example shows how to set a maximum of 2000 prefixes that can be redistributed into OSPFv3 process 1. If the number of prefixes redistributed reaches 75 percent of 2000 (1500 prefixes), a warning message is logged. Another warning is logged if the limit is reached and no more routes are redistributed.

```
RP/0/RSP0/CPU0:router(config)# router ospfv3 1
RP/0/RSP0/CPU0:router(config-ospfv3)# redistribute bgp 2406
RP/0/RSP0/CPU0:router(config-ospfv3)# maximum redistributed-prefixes 2000
```

mtu-ignore (OSPFv3)

To prevent the Open Shortest Path First Version 3 (OSPFv3) router process from checking whether neighbors are using the same maximum transmission unit (MTU) on a common interface when exchanging database descriptor (DBD) packets, use the **mtu-ignore** command in an appropriate configuration mode. To return to the default state, use the **no** form of this command.

mtu-ignore [**disable**]
no mtu-ignore

Syntax Description

disable	(Optional) Disables the attribute in instances in which it is specified at a higher level in the configuration.

Note
The **disable** keyword is not available in router ospfv3 configuration mode.

Command Default

If this command is not specified in interface configuration mode, then the interface adopts the MTU ignore parameter specified by the area.

If this command is not specified in area configuration mode, then the interface adopts the MTU ignore parameter specified for the process.

If this command is not specified at any level, then OSPFv3 checks the MTU received from neighbors when exchanging DBD packets.

Command Modes

Interface configuration

Area configuration

Router OSPFv3 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **mtu-ignore** command to check whether OSPFv3 neighbors are using the same MTU on a common interface. This check is performed when neighbors exchange DBD packets. If the receiving MTU in the DBD packet is higher than the MTU configured on the incoming interface, OSPF adjacency is not established.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to disable MTU mismatch detection on received DBD packets on GigabitEthernet interface 0/1/0/3:

```
RP/0/RSP0/CPU0:router(config)# router ospfv3 1
RP/0/RSP0/CPU0:router(config-ospfv3)# area 0
RP/0/RSP0/CPU0:router(config-ospfv3-ar)# interface GigabitEthernet 0/1/0/3
RP/0/RSP0/CPU0:router(config-ospfv3-ar-if)# mtu-ignore
```

neighbor (OSPFv3)

To configure Open Shortest Path First Version 3 (OSPFv3) routers interconnecting to nonbroadcast networks, use the **neighbor** command in interface configuration mode. To remove a configuration, use the **no** form of this command.

neighbor *ipv6-address* [**priority** *number*] [**poll-interval** *seconds*] [**cost** *number*] [**database-filter all out**]
no neighbor *ipv6-address* [**priority** *number*] [**poll-interval** *seconds*] [**cost** *number*] [**database-filter all out**]

Syntax Description

ipv6-address	Link- local IP Version 6 (IPv6) address of the neighbor. This argument must be in the form documented in RFC 2373, in which the address is specified in hexadecimal using 16-bit values between colons.
priority *number*	(Optional) Specifies an 8-bit number indicating the router priority value of the nonbroadcast neighbor associated with the IP address specified. The **priority** keyword does not apply to point-to-multipoint interfaces.
poll-interval *seconds*	(Optional) Specifies an unsigned integer value (in seconds) reflecting the poll interval. RFC 1247 recommends that this value be much larger than the hello interval. The **poll-interval** keyword does not apply to point-to-multipoint interfaces.
cost *number*	(Optional) Assigns a cost to the neighbor, in the form of an integer from 1 to 65535. Neighbors with no specific cost configured assume the cost of the interface, based on the **cost** command. On point-to-multipoint interfaces, **cost** *number* is the only keyword and argument combination that works. The **cost** keyword does not apply to nonbroadcast multiaccess (NBMA) networks.
database-filter all out	(Optional) Filters outgoing link-state advertisements (LSAs) to an OSPFv3 neighbor.

Command Default

No configuration is specified.

priority *number* : 0

poll-interval *seconds* : 120 seconds (2 minutes)

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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

One neighbor entry must be included in the software configuration for each known nonbroadcast network neighbor. The neighbor address must be the IPv6 link-local address of the interface.

If a neighboring router has become inactive (hello packets have not been seen for the router dead interval period), it may still be necessary to send hello packets to the dead neighbor. These hello packets are sent at a reduced rate called the poll interval.

When the router starts up, it sends only hello packets to those routers with nonzero priority; that is, routers that are eligible to become designated routers (DRs) and backup designated routers (BDRs). After the DR and BDR are selected, the DR and BDR then start sending hello packets to all neighbors to form adjacencies.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to declare a router at address fe80::3203:a0ff:fe9d:f3fe on a nonbroadcast network:

```
RP/0/RSP0/CPU0:router(config)# router ospfv3 1
RP/0/RSP0/CPU0:router(config-ospfv3)# area 0
RP/0/RSP0/CPU0:router(config-ospfv3-ar)# interface GigabitEthernet 0/2/0/0
RP/0/RSP0/CPU0:router(config-ospfv3-ar)# network non-broadcast
RP/0/RSP0/CPU0:router(config-ospfv3-ar-if)# neighbor fe80::3203:a0ff:fe9d:f3fe
```

Related Commands

Command	Description
priority (OSPFv3), on page 1354	Sets the router priority, which helps determine the designated router for this network.

network (OSPFv3)

To configure the Open Shortest Path First Version 3 (OSPFv3) network type to a type other than the default for a given medium, use the **network** command in an appropriate configuration mode. To return to the default value, use the **no** form of this command.

network {**broadcast** | **non-broadcast** | {**point-to-multipoint** [**non-broadcast**] | **point-to-point**}}
no network

Syntax Description

broadcast	Sets the network type to broadcast.
non-broadcast	Sets the network type to nonbroadcast multiaccess (NBMA).
point-to-multipoint	Sets the network type to point-to-multipoint.
[**non-broadcast**]	(Optional) Sets the point-to-multipoint network to be nonbroadcast. If you use the **non-broadcast** keyword, the **neighbor** command is required.
point-to-point	Sets the network type to point-to-point.

Command Default

If this command is not specified in interface configuration mode, then the interface adopts the network parameter specified by the area.

If this command is not specified in area configuration mode, then the interface adopts the network parameter specified for the process.

If this command is not specified at any level, then the OSPFv3 network type is the default of the given medium.

Command Modes

Interface configuration

Area configuration

Router OSPFv3 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **network** command to configure broadcast networks as NBMA networks when, for example, routers in your network do not support multicast addressing.

Most times, it is assumed that when you configure NBMA networks as either broadcast or nonbroadcast, there are virtual circuits from every router to every router or fully meshed network. However, there are other configurations where this assumption is not true; for example, a partially meshed network. In these cases, you can configure the OSPFv3 network type as a point-to-multipoint network. Routing between two routers that are not directly connected goes through the router that has virtual circuits to both routers. You need not configure neighbors when using this command.

If the **network** command is issued on an interface that does not allow it, this command is ignored.

OSPFv3 has two features related to point-to-multipoint networks. One feature applies to broadcast networks and the other feature applies to nonbroadcast networks:

- On point-to-multipoint, broadcast networks, you can use the **neighbor** command, and you must specify a cost to that neighbor.
- On point-to-multipoint, nonbroadcast networks, you must use the **neighbor** command to identify neighbors. Assigning a cost to a neighbor is optional.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to configure an Ethernet interface as point-to-point:

```
RP/0/RSP0/CPU0:router(config)# router ospfv3 1
RP/0/RSP0/CPU0:router(config-ospfv3)# area 0
RP/0/RSP0/CPU0:router(config-ospfv3-ar)# interface TenGigE0/1/0/3
RP/0/RSP0/CPU0:router(config-ospfv3-ar-if)# network point-to-point
```

Related Commands

Command	Description
neighbor (OSPFv3), on page 1342	Configures OSPFv3 routers interconnecting to nonbroadcast networks.

nssa (OSPFv3)

To configure an area as a not-so-stubby area (NSSA), use the **nssa** command in area configuration mode. To remove the NSSA distinction from the area, use the **no** form of this command.

nssa [**no-redistribution**] [**default-information-originate** [**metric** *metric-value* | **metric-type** *type-value*]] [**no-summary**]
no nssa

Syntax Description

no-redistribution	(Optional) Imports routes into the normal areas, but not into the NSSA area, by the **redistribute** command when the router is an NSSA area border router (ABR).
default-information-originate	(Optional) Generates a Type 7 default into the NSSA area. This keyword takes effect only on an NSSA ABR or NSSA autonomous system boundary router (ASBR).
metric *metric-value*	(Optional) Specifies a metric used for generating the default route. If you do not specify a default route metric value using the **nssa** and **defaultmetric** commands, the default metric value is 10. The value used is specific to the protocol.
metric-type *type-value*	(Optional) Specifies an external link type associated with the default route advertised into the Open Shortest Path First Version 3 (OSPFv3) routing domain. It can be one of the following values: **1**—Type 1 external route **2**—Type 2 external route
no-summary	(Optional) Prevents an (ABR) from sending summary link advertisements into the NSSA area.

Command Default

No NSSA area is defined.

If you do not specify a value using the **default-metric** command, the default metric value is 10.

The default *type-value* is Type 2 external route.

Command Modes

Area 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

A default route need not be defined in an NSSA ABR when the **nssa** command is configured. However, if this command is configured on an NSSA ASBR, then a default route must be defined.



Note NSSA cannot be configured for area 0 (backbone area).

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to configure area 1 as an NSSA area:

```
RP/0/RSP0/CPU0:router(config)# router ospfv3 1
RP/0/RSP0/CPU0:router(config-ospfv3)# router-id 10.18.1.1
RP/0/RSP0/CPU0:router(config-ospfv3)# area 1
RP/0/RSP0/CPU0:router(config-ospfv3-ar)# nssa
```

nsr (OSPFv3)

To configure nonstop routing (NSR) for the Open Shortest Path First Version 3 (OSPFv3) protocol, use the **nsr** command in OSPFv3 router configuration mode. To remove this command from the configuration file, use the **no** form of this command.

nsr [**disable**]
no nsr [**disable**]

Syntax Description

This command has no keywords or arguments.

Command Default

NSR is enabled.

Command Modes

OSPFv3 Router configuration

Command History

Release	Modification
Release 4.2.0	This command was introduced.
Release 6.0.0	This command was modified. NSR was enabled by default.

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.

OSPFv3 NSR will be disabled at process startup, by default. When enabled, this state is remembered in the active process, and, is regardless of the presence and pairing state of a standby RP, as well as the state of the standby process.

NSR can be enabled for multiple OSPFv3 processes. The maximum number of processes on which NSR can be enabled is four.

Task ID

Task ID	Operation
ospf	read, write

This example shows how to configure NSR for OSPFv3 process 211:

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#router ospfv3 211
RP/0/RSP0/CPU0:router(config-ospfv3)#nsr
```

This example shows how to configure NSR for OSPFv3 process 211:

```
RP/0/RSP0/CPU0:router#configure
```

```
RP/0/RSP0/CPU0:router(config)#router ospfv3 211
RP/0/RSP0/CPU0:router(config-ospfv3)#nsr disable
```

Related Commands

Command	Description
router ospfv3, on page 1367	
show ospfv3, on page 1370	

ospfv3 name-lookup

To configure Open Shortest Path First Version 3 (OSPFv3) to look up Domain Name System (DNS) names, use the **ospfv3 name-lookup** command in global configuration mode. To disable this function, use the **no** form of this command.

ospfv3 name-lookup
no ospfv3 name-lookup

Syntax Description

This command has no arguments or keywords.

Command Default

Routers are displayed by router ID or neighbor ID.

Command Modes

Global 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **ospfv3 name-lookup** command to simplify the task of searching for a router. Routers are displayed by name rather than by router ID or neighbor ID.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to configure OSPFv3 to look up DNS names for use in all OSPFv3 **show** command displays:

```
RP/0/RSP0/CPU0:router(config)# ospfv3 name-lookup
```

packet-size (OSPFv3)

To configure the size of Open Shortest Path First Version 3 (OSPFv3) packets up to the size specified by the maximum transmission unit (MTU), use the **packet-size** command in an appropriate configuration mode. To disable this function and reestablish the default packet size, use the **no** form of this command.

packet-size *bytes*
no packet-size

Syntax Description

bytes	Size in bytes. Range is 256 to 10000 bytes.

Command Default

If not specified, the default packet size is 1500 bytes.

Command Modes

Router OSPFv3 configuration

Area configuration

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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **packet-size** command to customize the size of OSPFv3 packets. The OSPFv3 protocol compares the packet size and the MTU size and uses the lower packet size value.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to configure the packet size:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router ospf osp3
RP/0/RSP0/CPU0:router(config-ospfv3)# packet-size 3500
```

passive (OSPFv3)

To suppress the sending of Open Shortest Path First Version 3 (OSPFv3) packets on an interface, use the **passive** command in an appropriate configuration mode. To remove the passive configuration, use the **no** form of this command.

passive [**disable**]
no passive

Syntax Description

disable	(Optional) Sends OSPFv3 updates.

Note
The **disable** keyword is not available in router ospfv3 configuration mode.

Command Default

If this command is not specified in interface configuration mode, then the interface adopts the passive parameter specified by the area.

If this command is not specified in area configuration mode, then the interface adopts the passive parameter specified for the process.

If this command is not specified at any level, then the passive parameter is disabled and OSPFv3 updates are sent on the interface.

Command Modes

Interface configuration

Area configuration

Router OSPFv3 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

OSPF routing information is neither sent nor received through the specified router interface. The specified interface address appears as a stub network in the OSPF domain.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows that OSPFv3 updates run over GigabitEthernet interface 0/3/0/0, 0/2/0/0, and 0/2/0/2. All other interfaces suppress sending OSPFv3 updates because they are in passive mode.

```
RP/0/RSP0/CPU0:router(config)# router ospfv3 1
RP/0/RSP0/CPU0:router(config-ospfv3)# router-id 10.0.0.206
RP/0/RSP0/CPU0:router(config-ospfv3)# passive
RP/0/RSP0/CPU0:router(config-ospfv3)# area 0
RP/0/RSP0/CPU0:router(config-ospfv3-ar)# interface GigabitEthernet 0/3/0/0
RP/0/RSP0/CPU0:router(config-ospfv3-ar-if)# passive disable
RP/0/RSP0/CPU0:router(config-ospfv3-ar-if)# exit
RP/0/RSP0/CPU0:router(config-ospfv3-ar)# interface GigabitEthernet 0/3/0/1
RP/0/RSP0/CPU0:router(config-ospfv3-ar-if)# exit
RP/0/RSP0/CPU0:router(config-ospfv3-ar)# exit
RP/0/RSP0/CPU0:router(config-ospfv3)# area 1
RP/0/RSP0/CPU0:router(config-ospfv3-ar)# passive disable
RP/0/RSP0/CPU0:router(config-ospfv3-ar)# interface GigabitEthernet 0/2/0/0
RP/0/RSP0/CPU0:router(config-ospfv3-ar-if)# exit
RP/0/RSP0/CPU0:router(config-ospfv3-ar)# interface GigabitEthernet 0/2/0/1
RP/0/RSP0/CPU0:router(config-ospfv3-ar-if)# passive
RP/0/RSP0/CPU0:router(config-ospfv3-ar-if)# exit
RP/0/RSP0/CPU0:router(config-ospfv3-ar)# interface GigabitEthernet 0/2/0/2
RP/0/RSP0/CPU0:router(config-ospfv3-ar-if)# exit
```

priority (OSPFv3)

To set the router priority for an interface, which helps determine the designated router for an Open Shortest Path First Version 3 (OSPFv3) link, use the **priority** command in an appropriate configuration mode. To return to the default value, use the **no** form of this command.

priority *value*
no priority

Syntax Description

value	8-bit unsigned integer indicating the router priority value. Range is 0 to 255.

Command Default

If this command is not specified in interface configuration mode, then the interface adopts the priority parameter specified by the area.

If this command is not specified in area configuration mode, then the interface adopts the priority parameter specified by the process.

If this command is not specified at any level, then the default priority is 1.

Command Modes

Interface configuration

Area configuration

Router OSPFv3 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

When two routers attached to a network both attempt to become the designated router, the one with the higher router priority takes precedence. If there is a tie, the router with the higher router ID takes precedence. A router with a router priority set to zero is ineligible to become the designated router or backup designated router. Router priority is configured only for interfaces to broadcast and nonbroadcast multiaccess (NBMA) networks.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to set the router priority value to 4 on GigabitEthernet interface 0/1/0/1:

```
RP/0/RSP0/CPU0:router(config)# router ospfv3 1
```

```
RP/0/RSP0/CPU0:router(config-ospfv3)# area 0
RP/0/RSP0/CPU0:router(config-ospfv3-ar)# interface GigabitEthernet 0/1/0/1
RP/0/RSP0/CPU0:router(config-ospfv3-ar-if)# priority 4
```

Related Commands

Command	Description
neighbor (OSPFv3), on page 1342	Configures OSPFv3 routers interconnecting to nonbroadcast networks.
network (OSPFv3), on page 1344	Configures the OSPFv3 network type to a type other than the default for a given medium.

protocol shutdown (OSPFv3)

To disable an instance of the Open Shortest Path First protocol, version 3 (OSPFv3), so that it cannot form an adjacency on any interface, use the **protocol shutdown** command in the ospfv3 configuration mode. To re-enable the OSPF protocol, use the **no** form of this command.

protocol shutdown
no protocol shutdown

Command Default

None

Command Modes

ospfv3 configuration

Command History

Release	Modification
Release 5.1	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.

Use the **protocol shutdown** command to disable the OSPFv3 protocol for a specific routing instance without removing any existing OSPF configuration parameters.

The OSPFv3 protocol continues to run on the router and you can use the current OSPFv3 configuration, but OSPFv3 does not form any adjacencies on any interface.

This command is similar to performing the **no router ospf** command.

Task ID

Task ID	Operations
ospf	read, write

Examples

This example shows how to disable the OSPFv3:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router ospfv3 100
RP/0/RSP0/CPU0:router(config-ospfv3)# protocol shutdown
```

range (OSPFv3)

To consolidate and summarize routes at an area boundary for Open Shortest Path First Version 3 (OSPFv3), use the **range** command in area configuration mode. To restore the default values, use the **no** form of this command.

range *ipv6-prefix/prefix-length* [**advertise** | **not-advertise**] [**cost** *number*]
no range *ipv6-prefix/prefix-length* [**advertise** | **not-advertise**] [**cost** *number*]

Syntax Description

ipv6-prefix	Summary prefix designated for a range of IP Version 6 (IPv6) prefixes. This argument must be in the form documented in RFC 2373, in which the address is specified in hexadecimal using 16-bit values between colons.
/ *prefix-length*	Length of the IPv6 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.
advertise	(Optional) Sets the address range status to advertise and generates a Type 3 summary link-state advertisement (LSA).
not-advertise	(Optional) Sets the address range status to DoNotAdvertise. The Type 3 summary LSA is suppressed and the component networks remain hidden from other networks.
cost *number*	(Optional) Specifies a cost for the range. Range is 1 to 16777214.

Command Default

Routes are not consolidated and summarized for an area.

Command Modes

Area 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **range** command only with Area Border Routers (ABRs). It is used to consolidate or summarize routes for an area. The result is that a single summary route is advertised to other areas by the ABR. Routing information is condensed at area boundaries. External to the area, a single route is advertised for each address range. This process is called *route summarization* .

You can use the **range** command to configure multiple ranges. Thus, OSPFv3 can summarize addresses for many different sets of address ranges.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to specify one summary route to be advertised by the ABR to other areas for all IPv6 prefixes within the range defined by summary prefix 4004:f000::/32:

```
RP/0/RSP0/CPU0:router(config)# router ospfv3 201
RP/0/RSP0/CPU0:router(config-ospfv3)# area 0
RP/0/RSP0/CPU0:router(config-ospfv3-ar)# range 4004:f000::/32
```

redistribute (OSPFv3)

To redistribute routes from one routing domain into Open Shortest Path First Version 3 (OSPFv3), use the **redistribute** command in an appropriate configuration mode. To remove the **redistribute** command from the configuration file and restore the system to its default condition in which the software does not redistribute routes, use the **no** form of this command.

Border Gateway Protocol (BGP)
redistribute bgp *process-id* [**metric** *metric-value*] [**metric-type** {**1** | **2**}] [**policy** *policy-name*] [**tag** *tag-value*]
no redistribute bgp *process-id* [**metric** *metric-value*] [**metric-type** {**1** | **2**}] [**policy** *policy-name*] [**tag** *tag-value*]

Local Interface Routes
redistribute connected [**metric** *metric-value*] [**metric-type** {**1** | **2**}] [**policy** *policy-name*] [**tag** *tag-value*]
no redistribute connected [**metric** *metric-value*] [**metric-type** {**1** | **2**}] [**policy** *policy-name*] [**tag** *tag-value*]

Enhanced Interior Gateway Routing Protocol (EIGRP)
redistribute eigrp *process-id* [**match** {**external** [**1** | **2**] | **internal**}] [**metric** *metric-value*] [**metric-type** {**1** | **2**}] [**route-policy** *policy-name*] [**tag** *tag-value*]
no redistribute eigrp *process-id* [**match** {**external** [**1** | **2**] | **internal**}] [**metric** *metric-value*] [**metric-type** {**1** | **2**}] [**route-policy** *policy-name*] [**tag** *tag-value*]

Intermediate System-to-Intermediate System (IS-IS)
redistribute isis *process-id* [**level-1** | **level-2** | **level-1-2**] [**metric** *metric-value*] [**metric-type** {**1** | **2**}] [**policy** *policy-name*] [**tag** *tag-value*]
no redistribute isis *process-id* [**level-1** | **level-2** | **level-1-2**] [**metric** *metric-value*] [**metric-type** {**1** | **2**}] [**policy** *policy-name*] [**tag** *tag-value*]

Open Shortest Path First Version 3 (OSPFv3)
redistribute ospfv3 *process-id* **match** {**external** | **1** | **2** | **internal** | **nssa-external** | [**1** | **2**]} [**metric** *metric-value*] [**metric-type** {**1** | **2**}] [**policy** *policy-name*] [**tag** *tag-value*]
no redistribute ospfv3 *process-id* [**match** | {**external** | **internal** | **nssa-external**}] [**metric** *metric-value*] [**metric-type** {**1** | **2**}] [**policy** *policy-name*] [**tag** *tag-value*]

Static
redistribute static [**metric** *metric-value*] [**metric-type** {**1** | **2**}] [**policy** *policy-name*] [**tag** *tag-value*]
no redistribute static [**metric** *metric-value*] [**metric-type** {**1** | **2**}] [**policy** *policy-name*] [**tag** *tag-value*]

Syntax Description

bgp	Distributes routes from the BGP protocol.

process-id	For the **bgp** keyword, an autonomous system number has the following ranges: • 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. For the **isis** keyword, an IS-IS instance name from which routes are to be redistributed. The value takes the form of a string. A decimal number can be entered, but it is stored internally as a string. For the **ospf** keyword, an OSPF instance name from which routes are to be redistributed. The value takes the form of a string. A decimal number can be entered, but it is stored internally as a string.
metric *metric-value*	(Optional) Specifies the metric used for the redistributed route. Range is 1 to 16777214. Use a value consistent with the destination protocol.
metric-type { **1** | **2** }	(Optional) Specifies the external link type associated with the route advertised into the OSPF routing domain. It can be one of two values: • **1**—Type 1 external route • **2**—Type 2 external route If no **metric-type** is specified, the default is Type 2 external routes.
tag *t ag-value*	(Optional) Specifies the 32-bit dotted-decimal value attached to each external route. This value is not used by the OSPF protocol itself, but is carried in the External LSAs. Range is 0 to 4294967295.
policy *policy-name*	(Optional) Specifies the identifier of a configured policy. A policy is used to filter the importation of routes from this source routing protocol to OSPF.
connected	Distributes routes that are established automatically by virtue of having enabled IP on an interface.
eigrp	Distributes routes from the EIGRP protocol.
isis	Distributes routes from the IS-IS protocol.
level-1	(Optional) Redistributes Level 1 routes into other IP routing protocols independently.
level-1- 2	(Optional) Redistributes both Level 1 and Level 2 routes into other IP routing protocols.
level-2	(Optional) Redistributes Level 2 routes into other IP routing protocols independently.
ospf	Distributes routes from the OSPF protocol.

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. If no match is specified, the default is no filtering.
static	Redistributes IP static routes.

Command Default

Route redistribution is disabled.

metric *metric-value:* Default is 20 for routes from all protocols except BGP routes, in which the default is 1.

metric-type *type-value* : Type 2 external route

All routes from the OSPFv3 routing protocol are redistributed.

tag *tag-value* : If no value is specified, the remote autonomous system number is used for routes from Border Gateway Protocol (BGP); for other protocols, the default is 0.

Command Modes

Router OSPFv3 configuration

Command History

Release	Modification
Release 3.7.2	This command was introduced.
Release 3.9.0	Asplain format for 4-byte Autonomous system numbers notation was supported.

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.

Note When redistributing routes (into OSPF) 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.

Disabling or changing the arguments of any keyword does not affect the state of other keywords.

In general, route redistribution from Level 1 to Level 2 is automatic. You might want to use this command to better control which Level 1 routes can be redistributed into Level 2.

The redistibution of Level 2 routes into Level 1 is called *route leaking*. Route leaking is disabled by default. That is, Level 2 routes are not automatically included in Level 1 link-state protocols. If you want to leak Level 2 routes into Level 1, you must enable that behavior by using this command.

Redistribution from Level 1 into Level 1 and from Level 2 into Level 2 is not allowed.

A router receiving a link-state packet with an internal metric considers the cost of the route from itself to the redistributing router plus the advertised cost to reach the destination. An external metric considers only the advertised metric to reach the destination.

Redistributed routing information should always be filtered by the **distribute-list prefix-list out** command. Use of this command ensures that only those routes intended by the administrator are passed along to the receiving routing protocol.

OSPFv3 Considerations

Whenever you use the **redistribute** or the **default-information** command to redistribute routes into an OSPFv3 routing domain, the router automatically becomes an ASBR. However, an ASBR does not, by default, generate a default route into the OSPFv3 routing domain.

When routes are redistributed between OSPFv3 processes, no OSPFv3 metrics are preserved.

When routes are redistributed into OSPF and no metric is specified with the **metric** keyword, OSPF uses 20 as the default metric for routes from all protocols except BGP routes, which get a metric of 1. Furthermore, when the router redistributes from one OSPFv3 process to another OSPFv3 process on the same router, and if no default metric is specified, the metrics in one process are carried to the redistributing process.

BGP Considerations

The only connected routes affected by this command are the routes not specified by the **network** (BGP) command.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to cause static routes to be redistributed into an OSPFv3 domain:

```
RP/0/RSP0/CPU0:router(config)# router ospfv3 109
RP/0/RSP0/CPU0:router(config-ospfv3)# redistribute isis level-1
```

Related Commands

Command	Description
default-information originate (OSPFv3), on page 1310	Redistributes routes into a routing domain.
distribute-list prefix-list out, on page 1320	Filters the routes redistributed into OSPFv3 from other routing protocols.

retransmit-interval (OSPFv3)

To specify the time between link-state advertisement (LSA) retransmissions for adjacencies belonging to the Open Shortest Path First Version 3 (OSPFv3) interface, use the **retransmit-interval** command in an appropriate configuration mode. To return to the default value, use the **no** form of this command.

retransmit-interval *seconds*
no retransmit-interval

Syntax Description

seconds	Time (in seconds) between retransmissions. It must be greater than the expected round-trip delay between any two routers on the attached network. Range is 1 to 65535 seconds.

Command Default

If this command is not specified in interface configuration mode, then the interface adopts the retransmit interval parameter specified by the area.

If this command is not specified in area configuration mode, then the interface adopts the retransmit interval parameter specified by the process.

If this command is not specified at any level, then the default retransmit interval is 5 seconds.

Command Modes

Interface configuration

Area configuration

Router OSPFv3 configuration

Virtual-link 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

When a router sends an LSA to its neighbor, it keeps the LSA until it receives the acknowledgment message. If the router receives no acknowledgment, it resends the LSA.

The setting of this parameter should be conservative, or needless retransmission results. The value should be larger for serial lines and virtual links.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to set the retransmit interval value to 8 seconds while in interface configuration mode:

```
RP/0/RSP0/CPU0:router(config)# router ospfv3 1
RP/0/RSP0/CPU0:router(config-ospfv3)# area 0
RP/0/RSP0/CPU0:router(config-ospfv3-ar)# interface GigabitEthernet 0/2/0/0
RP/0/RSP0/CPU0:router(config-ospfv3-ar-if)# retransmit-interval 8
```

router-id (OSPFv3)

To configure a router ID for the Open Shortest Path First Version 3 (OSPFv3) routing process, use the **router-id** command in an appropriate configuration mode. To cause the software to use the default method of determining the router ID, use the **no** form of this command after clearing or restarting the OSPF process.

router-id *router-id*
no router-id *router-id*

Syntax Description

router-id	32-bit router ID value specified in four-part, dotted-decimal notation.

Command Default

If this command is not configured, the router ID is the highest IP address for an interface on the router, with any loopback interface taking precedence.

Command Modes

Router OSPFv3 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

We recommend that you use the **router-id** command to explicitly specify a unique 32-bit numeric value for the router ID. This configuration ensures that OSPFv3 can function regardless of the interface address configuration. Clear the OSPF process using the **clear ospf process** command or restart the OSPF process for the **no router-id** command to take effect.

In router OSPFv3 configuration mode, OSPF attempts to obtain a router ID in the following ways (in order of preference):

1. By default, when the OSPF process initializes, it checks if there is a router-id in the checkpointing database.
2. The 32-bit numeric value specified by the OSPF **router-id** command in router configuration mode. (This value can be any 32-bit value. It is not restricted to the IPv4 addresses assigned to interfaces on this router, and need not be a routable IPv4 address.)
3. A global router ID provided by the system (possibly, the first loopback address found at the boot time.

If the OSPFv3 process cannot obtain a router ID from any of these sources, the router issues the following error message:

```
%OSPFv3-4-NORTRID : OSPFv3 process 1 cannot run - configure a router ID for this process
```

At this point, OSPFv3 is effectively passive on all its interfaces. To run OSPFv3, make a router ID available by one of the methods described.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to assign the IP address of 10.0.0.10 to the OSPFv3 process 109:

```
RP/0/RSP0/CPU0:router(config)# router ospfv3 109
RP/0/RSP0/CPU0:router(config-ospfv3)# router-id 10.0.0.10
```

Related Commands

Command	Description
clear ospfv3 process, on page 1296	Resets an OSPFv3 router process without stopping and restarting it.

router ospfv3

To configure an Open Shortest Path First Version 3 (OSPFv3) routing process, use the **router ospfv3** command in global configuration mode. To terminate an OSPFv3 routing process, use the **no** form of this command.

router ospfv3 *process-name*
no router ospfv3 *process-name*

Syntax Description

process-name	Name that uniquely identifies an OSPFv3 routing process. The process name is any alphanumeric string no longer than 40 characters.

Command Default

No OSPFv3 routing process is defined.

Command Modes

Global 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

You can specify multiple OSPFv3 routing processes in each router. Up to 10 processes can be configured. The recommendation is not to exceed 4 OSPFv3 processes.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to instantiate an OSPFv3 routing process with a process name of 1:

```
RP/0/RSP0/CPU0:router(config)# router ospfv3 1
```

Related Commands

Command	Description
area (OSPFv3), on page 1289	Defines an OSPFv3 area.
interface (OSPFv3), on page 1333	Defines an OSPFv3 interface by type.

sham-link (OSPFv3)

To configure an Open Shortest Path First version 3 (OSPFv3) sham link between two provider edge routers on a non-default VRF, use the **sham-link** command in OSPFv3 router area sub-configuration mode. To terminate an OSPFv3 sham link, use the **no** form of this command.

sham-link *source-address destination-address*
no sham-link

Syntax Description

source-address	IP address of the local (source) sham-link endpoint specified in four-part, dotted-decimal notation.
destination-address	IP address of the remote (destination) sham-link endpoint specified in four-part, dotted-decimal notation.

Command Default

No sham link is configured.

Command Modes

OSPFv3 router area sub-configuration.

Command History

Release	Modification
Release 5.1	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.

Use the **sham-link** command to configure a point-to-point connection between two provider edge (PE) routers creating an interconnect between two VPN sites (VPN backbone). Sham links are configured on PE provider edge (PE) routers in a Multiprotocol Label Switching (MPLS) VPN backbone.

Task ID

Task ID	Operations
ospf	read, write

Examples

This example shows how to configure an OSPFv3 sham link:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router ospfv3 ospfv3
RP/0/RSP0/CPU0:router(config_ospfv3)# vrf 1
RP/0/RSP0/CPU0:router(config_ospfv3_vrf)# area 1
RP/0/RSP0/CPU0:router(config_ospfv3_vrf_ar)# sham-link 100::1 200::1
RP/0/RSP0/CPU0:router(config_ospf_vrf_ar_sl)# cost 23
```

Related Commands

Command	Description
area (OSPFv3), on page 1289	Configures an OSPF area.
cost (OSPFv3), on page 1302	Explicitly specifies the cost of the interface (network) for OSPF path calculation.

show ospfv3

To display general information about Open Shortest Path First Version 3 (OSPFv3) routing processes, use the **show ospfv3** command in EXEC mode.

show ospfv3 [*process-name*]

Syntax Description

process-name	(Optional) Name that uniquely identifies an OSPFv3 routing process. The process name is defined by the **router ospfv3** command. If this argument is included, only information for the specified routing process is displayed.

Command Default

None

Command Modes

EXEC

Command History

Release	Modification
Release 3.7.2	This command was introduced.
Release 3.9.0	Asplain format for 4-byte Autonomous system numbers notation was supported. The input parameters and output were modified to display 4-byte autonomous system numbers and extended communities in either asplain or asdot notations.

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
ospf	read

Examples

The following is a sample output from the **show ospfv3** command:

```
RP/0/RSP0/CPU0:router# show ospfv3 1

Routing Process "ospfv3 test" with ID 3.3.3.3
It is an autonomous system boundary router
Redistributing External Routes from,
static
 Maximum number of redistributed prefixes 10240
 Threshold for warning message 75%
Initial SPF schedule delay 5000 msecs
Minimum hold time between two consecutive SPFs 10000 msecs
Maximum wait time between two consecutive SPFs 10000 msecs
Initial LSA throttle delay 0 msecs
Minimum hold time for LSA throttle 5000 msecs
Maximum wait time for LSA throttle 5000 msecs
Minimum LSA arrival 1000 msecs
LSA group pacing timer 240 secs
```

```
Interface flood pacing timer 33 msecs
Retransmission pacing timer 66 msecs
Maximum number of configured interfaces 255
Number of external LSA 1. Checksum Sum 0x004468
Number of areas in this router is 1. 1 normal 0 stub 0 nssa
 Area BACKBONE(0) (Inactive)
  Number of interfaces in this area is 1
  SPF algorithm executed 1 times
  Number of LSA 3. Checksum Sum 0x018109
  Number of DCbitless LSA 0
  Number of indication LSA 0
  Number of DoNotAge LSA 0
  Flood list length 0
```

This table describes the significant fields shown in the display.

Table 119: show ospfv3 Field Descriptions

Field	Description
Routing Process “ospfv3 test” with ID	OSPFv3 process name.
It is	Types are internal, area border, or autonomous system boundary.
Redistributing External Routes from	Lists of redistributed routes, by protocol.
Maximum number of redistributed prefixes	Number of redistributed prefixes.
Threshold for warning message	Warning message threshold.
Initial SPF schedule delay	Delay time of SPF calculations.
Minimum hold time between two consecutive SPFs	Minimum hold time between consecutive SPFs.
Maximum wait time between two consecutive SPFs	Maximum wait time between consecutive SPFs.
Initial LSA throttle delay	Delay time of LSA throttle.
Maximum hold time for LSA throttle	After initial throttle delay, the LSA generation is backed off by hold interval.
Maximum wait time for LSA throttle	Maximum throttle delay for LSA generation.
Minimum LSA arrival	Minimum LSA arrival.
LSA group pacing timer	Configured LSA group pacing timer (in seconds).
Interface flood pacing timer	Flooding pacing interval.
Retransmission pacing timer	Retransmission pacing interval.
Maximum number of configured interfaces	Maximum number of configured interfaces.
Number of external LSA	Number of external LSAs.
Number of areas in this router is	Number of areas configured for the router.

Field	Description
Number of interfaces in this area is	Number of interfaces in the area.
SPF algorithm executed *n* times	Times SPF algorithm was executed.
Number of LSA	Number of LSAs.
Number of DCbitless LSA	Number of DCbitless LSAs.
Number of indication LSA	Number of indication LSAs.
Number of DoNotAge LSA	Number of do-not-age LSAs.
Flood list length	Flood list length.

Related Commands

Command	Description
router ospfv3, on page 1367	Configures an OSPFv3 routing process.

show ospfv3 border-routers

To display the internal Open Shortest Path First Version 3 (OSPFv3) routing table entries to an area border router (ABR) and autonomous system boundary router (ASBR), use the **show ospfv3 border-routers** command in EXEC mode.

show ospfv3 [*process-name*] **border-routers** [*router-id*]

Syntax Description

process-name	(Optional) Name that uniquely identifies an OSPFv3 routing process. The process name is defined by the **router ospfv3** command. If this argument is included, only information for the specified routing process is displayed.
router-id	(Optional) 32-bit router ID value specified in four-part, dotted-decimal notation.

Command Default

No default behavior or values

Command Modes

EXEC

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 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
ospf	read

Examples

The following is sample output from the **show ospfv3 border-routers** command:

```
RP/0/RSP0/CPU0:router# show ospfv3 border-routers

  OSPFv3 1 Internal Routing Table
  Codes: i - Intra-area route, I - Inter-area route
  i 10.0.0.207 [1] via fe80::3034:30ff:fe33:3742, GigabitEthernet 0/3/0/0, ABR/ASBR, Area
1, SPF 3
  i 10.0.0.207 [10] via fe80::204:c0ff:fe22:73fe, Ethernet0/0/0/0, ABR/ASBR, Area 0, SPF 7
```

This table describes the significant fields shown in the display.

Table 120: show ospf border-routers Field Descriptions

Field	Description
i	Type of this route; i indicates an intra-area route, I an inter-area route.

Field	Description
10.0.0.207	Router ID of destination.
[1]	Cost of using this route.
fe80::3034:30ff:fe33:3742	Next-hop toward the destination.
GigabitEthernet 0/3/0/0	Packets destined for fe80::3034:30ff:fe33:3742 are sent over GigabitEthernet interface 3/0/0/0.
ABR/ASBR	Router type of the destination; it is either an area border router (ABR) or autonomous system boundary router (ASBR) or both.
Area 1	Area ID of the area from which this route is learned.
SPF 3	Internal number of the shortest path first (SPF) calculation that installs this route.

Related Commands

Command	Description
router ospfv3, on page 1367	Configures an OSPFv3 routing process.

show ospfv3 database

To display lists of information related to the Open Shortest Path First Version 3 (OSPFv3) database for a specific router, use the **show ospfv3 database** command in EXEC mode.

show ospfv3 [*process-name* [*area-id*]] **database**
show ospfv3 [*process-name* [*area-id*]] **database**[**adv-router** [*router-id*]]
show ospfv3 [*process-name* [*area-id*]] **database** [**database-summary**]
show ospfv3 [*process-name* [*area-id*]] **database** [**external**] [*link-state-id*]
show ospfv3 [*process-name* [*area-id*]] **database** [**external**] [*link-state-id*] [**internal**] [**adv-router** [*router-id*]]
show ospfv3 [*process-name* [*area-id*]] **database** [**external**] [*link-state-id*] [**internal**] [**self-originate**]
show ospfv3 [*process-name* [*area-id*]] **database** [**grace**] [*link-state-id*] [**adv-router** [*router-id*]] [**internal**] [**self-originate**]
show ospfv3 [*process-name* [*area-id*]] **database** [**inter-area prefix**] [*link-state-id*]
show ospfv3 [*process-name* [*area-id*]] **database** [*link-state-id*] [**internal**] [**adv-router** [*router-id*]]
show ospfv3 [*process-name* [*area-id*]] **database** [**inter-area prefix**] [*link-state-id*] [**internal**] [**self-originate**]
show ospfv3 [*process-name* [*area-id*]] **database** [**inter-area router**] [*link-state-id*]
show ospfv3 [*process-name* [*area-id*]] **database** [**inter-area router**] [*link-state-id*] [**internal**] [**adv-router** [*router-id*]]
show ospfv3 [*process-name* [*area-id*]] **database** [**inter-area router**] [*link-state-id*] [**internal**] [**self-originate**]
show ospfv3 [*process-name* [*area-id*]] **database** [**link**] [*link-state-id*]
show ospfv3 [*process-name* [*area-id*]] **database** [**link**] [*link-state-id*] [**internal**] [**adv-router** [*router-id*]]
show ospfv3 [*process-name* [*area-id*]] **database** [**link**] [*link-state-id*] [**internal**] [**self-originate**]
show ospfv3 [*process-name* [*area-id*]] **database** [**network**] [*link-state-id*]
show ospfv3 [*process-name* [*area-id*]] **database** [**network**] [*link-state-id*] [**internal**] [**adv-router** [*router-id*]]
show ospfv3 [*process-name* [*area-id*]] **database** [**network**] [*link-state-id*] [**internal**] [**self-originate**]
show ospfv3 [*process-name* [*area-id*]] **database** [**nssa-external**] [*link-state-id*]
show ospfv3 [*process-name* [*area-id*]] **database** [**nssa-external**] [*link-state-id*] [**internal**] [**adv-router** [*router-id*]]
show ospfv3 [*process-name* [*area-id*]] **database** [**nssa-external**] [*link-state-id*] [**internal**] [**self-originate**]
show ospfv3 [*process-name* [*area-id*]] **database** [**prefix**] [**ref-lsa**] [**router** | **network**] [*link-state-id*] [**internal**] [**adv-router** [*router-id*]]
show ospfv3 [*process-name* [*area-id*]] **database** [**prefix**] [**ref-lsa**] [**router** | **network**] [*link-state-id*] [**internal**] [**self-originate**]
show ospfv3 [*process-name* [*area-id*]] **database** [**prefix**] [*link-state-id*]
show ospfv3 [*process-name* [*area-id*]] **database** [**prefix**] [*link-state-id*] [**internal**] [**adv-router** [*router-id*]]
show ospfv3 [*process-name* [*area-id*]] **database** [**prefix**] [*link-state-id*] [**internal**] [**self-originate**]
show ospfv3 [*process-name* [*area-id*]] **database** [**router**] [*link-state-id*]
show ospfv3 [*process-name area-id*] **database [router]** [**adv-router** [*router-id*]]
show ospfv3 [*process-name* [*area-id*]] **database** [**router**] [*link-state-id*] [**internal**] [**self-originate**]
show ospfv3 [*process-name* [*area-id*]] **database** [**self-originate**]

show ospfv3 [*process-name* [*area-id*]] **database** [**unknown** [**area** | **as** | **link**]] [*link-state-id*] [**internal**] [**adv-router** [*router-id*]] [**self-originate**]

Syntax Description

process-name	(Optional) Name that uniquely identifies an OSPFv3 routing process. The process name is defined by the **router ospfv3** command. If this argument is included, only information for the specified routing process is displayed.
area-id	(Optional) Area number used to define the particular area.
adv-router [*router-id*]	(Optional) Displays all link-state advertisements (LSAs) of the specified router.
asbr-summary	(Optional) Displays information only about the Autonomous System Boundary Router (ASBR) summary LSAs.
database-summary	(Optional) Displays how many of each type of LSA are in the database for each area and the total.
external	(Optional) Displays information only about external LSAs.
grace	(Optional) Displays information about the state for the graceful restart link.
internal	(Optional) Displays information only about internal LSAs.
self-originate	(Optional) Displays only self-originated LSAs (from the local router).
link-state-id	(Optional) LSA ID that uniquely identifies the LSA. For network LSAs and link LSAs, this ID is the interface ID of the link of the router originating the LSA.
inter-area prefix	(Optional) Displays information only about the interarea prefix LSAs.
inter-area router	(Optional) Displays information only about the interarea router LSAs.
link	(Optional) Displays information only about the link LSAs.
network	(Optional) Displays information only about the network LSAs.
nssa-external	(Optional) Displays information only about the not-so-stubby area (NSSA) external LSAs.
prefix	(Optional) Displays information only about the prefix LSAs.
ref-lsa	(Optional) Displays referenced LSA information.
router	(Optional) Displays information only about the router LSAs.
unknown	(Optional) Displays information only about unknown LSAs.
area	(Optional) Displays information only about the area LSAs.
as	(Optional) Displays information only about the autonomous system LSAs.

Command Default

No default behavior or values

Command Modes

EXEC

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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

The various forms of this command deliver information about different OSPFv3 link-state advertisements.

Task ID

Task ID	Operations
ospf	read

Examples

The following is sample output from the **show ospfv3 database** command with no arguments or keywords:

```
RP/0/RSP0/CPU0:router# show ospfv3 database

             OSPFv3 Router with ID (10.0.0.207) (Process ID 1)

                 Router Link States (Area 0)

  ADV Router      Age         Seq#       Fragment ID  Link count  Bits
  0.0.0.1         163         0x80000039 0            2           None
  10.0.0.206      145         0x80000005 0            1           EB
  10.0.0.207      151         0x80000004 0            1           EB
  192.168.0.0     163         0x80000039 0            1           None

                 Net Link States (Area 0)

  ADV Router      Age         Seq#       Link ID    Rtr count
  10.0.0.207      152         0x80000002 1          3
  192.168.0.0     163         0x80000039 1          2

                 Inter Area Prefix Link States (Area 0)

  ADV Router      Age         Seq#       Prefix
  10.0.0.206      195         0x80000001 3002::/56
  10.0.0.207      197         0x80000001 3002::/56
  10.0.0.206      195         0x80000001 3002::206/128
  10.0.0.207      182         0x80000001 3002::206/128

                 Inter Area Router Link States (Area 0)

  ADV Router      Age         Seq#       Link ID    Dest RtrID
  10.0.0.207      182         0x80000001 167772366  10.0.0.206
  10.0.0.206      182         0x80000001 167772367  10.0.0.207

                 Link (Type-8) Link States (Area 0)

  ADV Router      Age         Seq#       Link ID    Interface
  0.0.0.1         163         0x80000039 1          Et0/0/0/0
  10.0.0.207      202         0x80000001 1          Et0/0/0/0
  10.0.0.206      200         0x80000001 2          Et0/0/0/0

                 Intra Area Prefix Link States (Area 0)
```

```
ADV Router      Age         Seq#       Link ID    Ref-lstype  Ref-LSID
192.168.0.0     163         0x80000039 0          0x2002      1
192.168.0.0     163         0x80000039 1          0x2001      0
10.0.0.207      157         0x80000001 1001       0x2002      1
```

This table describes the significant fields shown in the display.

Table 121: show ospfv3 database Field Descriptions

Field	Description
ADV Router	ID of advertising router.
Age	Link-state age.
Seq#	Link-state sequence number (detects old or duplicate LSAs).
Fragment ID	Router LSA fragment ID.
Link count	Number of links described.
Bits	B indicates that the router is an area border router. E indicates that the router is an autonomous system boundary router. V indicates that the router is a virtual link endpoint. W indicates that the router is a wildcard multicast receiver.
Link ID	Unique LSA ID.
Rtr count	Number of routers attached to the link.
Prefix	Prefix of the route being described.
Dest RtrID	Router ID of the router being described.
Interface	Link described by the LSA.
Ref-lstype	LSA type of the LSA being referenced.
Ref-LSID	LSA ID of the LSA being referenced.

The following is sample output from the **show ospfv3 database** command with the **external** keyword:

```
RP/0/RSP0/CPU0:router# show ospfv3 database external

              OSPFv3 Router with ID (10.0.0.206) (Process ID 1)

                  Type-5 AS External Link States

    LS age: 189
    LS Type: AS External Link
    Link State ID: 0
    Advertising Router: 10.0.0.206
    LS Seq Number: 80000002
    Checksum: 0xa303
```

```
    Length: 36
    Prefix Address: 2222::
    Prefix Length: 56, Options: None
    Metric Type: 2 (Larger than any link state path)
    Metric: 20
    External Route Tag: 0
```

This table describes the significant fields shown in the display.

Table 122: show ospfv3 database external Field Descriptions

Field	Description
OSPFv3 Router with ID	Router ID number.
Process ID	OSPFv3 process name.
LS age	Link-state age.
LS Type	Link-state type.
Link State ID	Link-state ID.
Advertising Router	ID of Advertising router.
LS Seq Number	Link-state sequence number (detects old or duplicate LSAs).
Checksum	LS checksum (Fletcher checksum of the complete contents of the LSA).
Length	Length (in bytes) of the LSA.
Prefix Address	IPv6 address prefix of the route being described.
Prefix Length	Length of the IPv6 address prefix.
Metric Type	External type.
Metric	Link-state metric.
External Route Tag	External route tag, a 32-bit field attached to each external route. This tag is not used by the OSPFv3 protocol itself.

The following is sample output from the **show ospfv3 database** command with the **inter-area prefix** keyword:

```
RP/0/RSP0/CPU0:router# show ospfv3 database inter-area prefix

             OSPFv3 Router with ID (10.0.0.206) (Process ID 1)

                 Inter Area Prefix Link States (Area 0)

    LS age: 715
    LS Type: Inter Area Prefix Links
    Link State ID: 0
    Advertising Router: 10.0.0.206
    LS Seq Number: 80000002
    Checksum: 0x3cb5
```

```
    Length: 36
    Metric: 1
    Prefix Address: 3002::
    Prefix Length: 56, Options: None
```

This table describes the significant fields shown in the display.

Table 123: show ospfv3 database inter-area prefix Field Descriptions

Field	Description
OSPFv3 Router with ID	Router ID number.
Process ID	OSPFv3 process name.
LS age	Link-state age.
LS Type	Link-state type.
Link State ID	Link-state ID.
Advertising Router	ID of advertising router.
LS Seq Number	Link-state sequence (detects old or duplicate LSAs).
Checksum	Link-state checksum (Fletcher checksum of the complete contents of the LSA).
Length	Length (in bytes) of the LSA.
Metric	Link-state metric.
Prefix Address	IPv6 prefix of the route being described.
Prefix Length	IPv6 prefix length of the route being described.
Options	LA indicates that the prefix is a local address. MC indicates the prefix is multicast capable. NU indicates that the prefix is not unicast capable. P indicates that the prefix should be propagated at a not-so-stubby area (NSSA) area border.

The following is sample output from the **show ospfv3 database** command with the **inter-area router** keyword:

```
RP/0/RSP0/CPU0:router# show ospfv3 database inter-area router

              OSPFv3 Router with ID (10.0.0.206) (Process ID 1)

                  Inter Area Router Link States (Area 0)

    LS age: 1522
    Options: (V6-Bit E-Bit R-bit DC-Bit)
    LS Type: Inter Area Router Links
    Link State ID: 167772366
    Advertising Router: 10.0.0.207
    LS Seq Number: 80000002
    Checksum: 0xcaae
    Length: 32
    Metric: 1
```

```
    Destination Router ID: 10.0.0.206
```

This table describes the significant fields shown in the display.

Table 124: show ospfv3 database inter-area routerField Descriptions

Field	Description
OSPFv3 Router with ID	Router ID number.
Process ID	OSPFv3 process name.
LS age	Link-state age.
Options	Type of service options (Type 0 only): DC—Supports demand circuits.E—Capable of processing external LSAs. MC—Forwards IP multicast. N—Supports Type 7 LSAs. R—Router is active. V6—Include in IPv6 routing calculations.
LS Type	Link-state type.
Link State ID	Link-state ID.
Advertising Router	ID of the advertising router.
LS Seq Number	Link-state sequence (detects old or duplicate LSAs).
Checksum	Link-state checksum (Fletcher checksum of the complete contents of the LSA.)
Length	Length (in bytes) of the LSAs.
Metric	Link-state metric.
Destination Router ID	Router ID of the router being described.

The following is sample output from the **show ospfv3 database** command with the **link** keyword:

```
RP/0/RSP0/CPU0:router# show ospfv3 database link

              OSPFv3 Router with ID (10.0.0.206) (Process ID 1)

                  Link (Type-8) Link States (Area 0)

    LS age: 620
    Options: (V6-Bit E-Bit R-bit DC-Bit)
    LS Type: Link-LSA (Interface: Ethernet0/0/0/0)
    Link State ID: 1 (Interface ID)
    Advertising Router: 10.0.0.207
    LS Seq Number: 80000003
    Checksum: 0x7235
    Length: 56
    Router Priority: 1
    Link Local Address: fe80::204:c0ff:fe22:73fe
    Number of Prefixes: 1
    Prefix Address: 7002::
    Prefix Length: 56, Options: None
```

This table describes the significant fields shown in the display.

Table 125: show ospfv3 database link Field Descriptions

Field	Description
OSPFv3 Router with ID	Router ID number.
Process ID	OSPFv3 process name.
LS age	Link-state age.
Options	Type of service options (Type 0 only): DC—Supports demand circuits.E—Capable of processing external LSAs.MC—Forwards IP multicast.N—Supports type-7 LSAs.R—Router is active.V6—Include in IPv6 routing calculations.
LS Type	Link-state type.
Link State ID	Link-state ID (Interface ID).
Advertising Router	ID of the advertising router.
LS Seq Number	Link-state sequence (detects old or duplicate LSAs).
Checksum	Link-state checksum (Fletcher checksum of the complete contents of the LSA).
Length	Length (in bytes) of the LSAs.
Router Priority	Interface priority of originating router.
Link Local Address	Link local address of the interface.
Number of Prefixes	Number of prefixes associated with the link.
Prefix Address and Length	List of prefixes associated with the link.
Options	LA indicates that the prefix is a local address. MC indicates that the prefix is multicast capable. NU indicates that the prefix is not unicast capable. P indicates that the prefix should be propagated at an NSSA area border.

The following is sample output from the **show ospfv3 database** command with the **network** keyword:

```
RP/0/RSP0/CPU0:router# show ospfv3 database network

              OSPFv3 Router with ID (10.0.0.206) (Process ID 1)

                  Net Link States (Area 0)

    LS age: 1915
    Options: (V6-Bit E-Bit R-bit DC-Bit)
    LS Type: Network Links
    Link State ID: 1 (Interface ID of Designated Router)
```

```
    Advertising Router: 10.0.0.207
    LS Seq Number: 80000004
    Checksum: 0x4330
    Length: 36
          Attached Router: 10.0.0.207
          Attached Router: 0.0.0.1
          Attached Router: 10.0.0.206
```

This table describes the significant fields shown in the display.

Table 126: show ospfv3 database network Field Descriptions

Field	Description
OSPFv3 Router with ID	Router ID number.
Process ID 1	OSPFv3 process name.
LS age	Link-state age.
Options	Type of service options (Type 0 only): DC—Supports demand circuits.E—Capable of processing external LSAs.MC—Forwards IP multicast.N—Supports Type 7 LSAs.R—Router is active.V6—Include in IPv6 routing calculations.
LS Type	Link-state type.
Link State ID	Link-state ID of the designated router.
Advertising Router	ID of the advertising router.
LS Seq Number	Link-state sequence (detects old or duplicate LSAs).
Checksum	Link-state checksum (Fletcher checksum of the complete contents of the LSA).
Length	Length (in bytes) of the LSA.
Attached Router	List of routers attached to the network, by router ID.

The following is sample output from the **show ospfv3 database** command with the **prefix** keyword:

```
RP/0/RSP0/CPU0:router# show ospfv3 database prefix

             OSPFv3 Router with ID (10.0.0.206) (Process ID 1)

                 Intra Area Prefix Link States (Area 1)

    Routing Bit Set on this LSA
    LS age: 356
    LS Type: Intra-Area-Prefix-LSA
    Link State ID: 0
    Advertising Router: 10.0.0.206
    LS Seq Number: 8000001e
    Checksum: 0xcdaa
    Length: 44
    Referenced LSA Type: 2001
```

```
Referenced Link State ID: 0
Referenced Advertising Router: 10.0.0.206
Number of Prefixes: 1
Prefix Address: 8006::
Prefix Length: 56, Options: None, Metric: 1
```

This table describes the significant fields shown in the display.

Table 127: show ospfv3 database prefix Field Descriptions

Field	Description
OSPFv3 Router with ID	Router ID number.
Process ID 1	OSPFv3 process name.
LS age	Link-state age.
LS Type	Link-state type.
Link State ID	Link-state ID of the designated router.
Advertising Router	ID of the advertising router.
LS Seq Number	Link-state sequence (detects old or duplicate LSAs).
Checksum	Link-state checksum (Fletcher checksum of the complete contents of the LSA).
Length	Length (in bytes) of the LSA.
Referenced LSA Type	Router LSA or network LSA of the prefixes referenced.
Referenced Link State ID	Link-state ID of the router or network LSA.
Referenced Advertising Router	Advertising router of the referenced LSA.
Number of Prefixes	Number of prefixes listed in the LSA.
Prefix Address	Prefix associated with the router or network.
Prefix Length	Length of the prefix.
Options	LA indicates that the prefix is a local address. MC indicates that the prefix is multicast capable. NU indicates that the prefix is not unicast capable. P indicates the prefix should be propagated at an NSSA area border.
Metric	Cost of the prefix.

The following is sample output from the **show ospfv3 database** command with the **router** keyword:

```
RP/0/RSP0/CPU0:router# show ospfv3 database router

            OSPFv3 Router with ID (10.0.0.206) (Process ID 1)
```

```
                Router Link States (Area 0)

LS age: 814
Options: (V6-Bit E-Bit R-bit)
LS Type: Router Links
Link State ID: 0
Advertising Router: 0.0.0.1
LS Seq Number: 8000003c
Checksum: 0x51ca
Length: 56
Number of Links: 2

  Link connected to: a Transit Network
  Link Metric: 10
  Local Interface ID: 1
    Neighbor (DR) Interface ID: 1
    Neighbor (DR) Router ID: 10.0.0.207

  Link connected to: a Transit Network
  Link Metric: 10
  Local Interface ID: 2
    Neighbor (DR) Interface ID: 1
    Neighbor (DR) Router ID: 10.0.0.0
```

This table describes the significant fields shown in the display.

Table 128: show ospfv3 database router Field Descriptions

Field	Description
OSPFv3 Router with ID	Router ID number.
Process ID 1	OSPFv3 process name.
LS age	Link-state age.
Options	Type of service options (Type 0 only): DC—Supports demand circuits.E—Capable of processing external LSAs.MC—Forwards IP multicast. N—Supports Type 7 LSAs.R—Router is active.V6—Include in IPv6 routing calculations.
LS Type	Link-state type.
Link State ID	Link-state ID of the designated router.
Advertising Router	ID of the advertising router.
LS Seq Number	Link-state sequence (detects old or duplicate LSAs).
Checksum	Link-state checksum (Fletcher checksum of the complete contents of the LSA).
Length	Length (in bytes) of the LSA.

Field	Description
Link connected to	The type of network to which this interface is connected. Values are: • Another Router (point-to-point). • A Transit Network. • A Virtual Link.
Link Metric	OSPF cost of using this link.
Local Interface ID	Number that uniquely identifies an interface on a router.

Related Commands

Command	Description
router ospfv3, on page 1367	Configures an OSPFv3 routing process.

show ospfv3 flood-list

To display a list of Open Shortest Path First Version 3 (OSPFv3) link-state advertisements (LSAs) waiting to be flooded over an interface, use the **show ospfv3 flood-list** command in EXEC mode.

show ospfv3 [*process-name*] [*area-id*] **flood-list** [*type interface-path-id*]

Syntax Description

process-name	(Optional) Name that uniquely identifies an OSPFv3 routing process. The process name is defined by the **router ospfv3** command. If this argument is included, only information for the specified routing process is displayed.
area-id	(Optional) Area number used to define the particular area.
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

No default behavior or values

Command Modes

EXEC

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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **show ospfv3 flood-list** command to display OSPFv3 packet pacing.

Task ID

Task ID	Operations
ospf	read

Examples

The following sample output from the **show ospfv3 flood-list** command shows three entries for the OSPFv3 1 process running over GigabitEthernet interface 0/3/0/0:

```
RP/0/RSP0/CPU0:router# show ospfv3 flood-list GigabitEthernet 0/3/0/0

  Flood Lists for OSPFv3 1
```

```
Interface GigabitEthernet 0/3/0/0, Queue length 3
 Link state retransmission due in 24 msec

Displaying 3 entries from flood list:

 Type       LS ID           ADV RTR     Seq NO        Age  Checksum
    3        0.0.0.199      10.0.0.207 0x80000002    3600 0x00c924
    3        0.0.0.200      10.0.0.207 0x80000002    3600 0x008966
    4       10.0.0.206      10.0.0.207 0x80000008       0 0x001951
```

This table describes the significant fields shown in the display.

Table 129: show ospfv3 flood-list Field Descriptions

Field	Description
Interface	Interface for which information is displayed.
Queue length	Number of LSAs waiting to be flooded.
Link state retransmission due in	Length of time before next link-state transmission.
Type	Type of LSA.
LS ID	Link-state ID of the LSA.
ADV RTR	IP address of advertising router.
Seq NO	Sequence number of LSA.
Age	Age of LSA (in seconds).
Checksum	Checksum of LSA.

Related Commands

Command	Description
router ospfv3, on page 1367	Configures an OSPFv3 routing process.

show ospfv3 interface

To display Open Shortest Path First Version 3 (OSPFv3) interface information, use the **show ospfv3 interface** command in EXEC mode.

show ospfv3 [*process-name*] [*area-id*] **interface** [*type interface-path-id*]

Syntax Description

process-name	(Optional) Name that uniquely identifies an OSPFv3 routing process. The process name is defined by the **router ospfv3** command. If this argument is included, only information for the specified routing process is displayed.
area-id	(Optional) Area number used to define the particular area.
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

No default behavior or values

Command Modes

EXEC

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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **show ospfv3 interface** command when the adjacency between two neighboring routers is not forming. Adjacency means that the routers synchronize their databases when they discover each other.

You can look at the output to check the physical link and line protocol status and to confirm that the network type and timer intervals match those of the neighboring routers.

Task ID

Task ID	Operations
ospf	read

Examples

The following is sample output from the **show ospfv3 interface** command when GigabitEthernet interface 0/2 /0/0 is specified:

```
RP/0/RSP0/CPU0:router# show ospfv3 interface GigabitEthernet 0/
          2


             /

          0

          /0GigabitEthernet 0/2/0/0 is up, line protocol is up
  Link Local address fe80::203:a0ff:fe9d:f3fe, Interface ID 2
  Area 0, Process ID 1, Instance ID 0, Router ID 10.0.0.206
  Network Type BROADCAST, Cost: 10
  BFD enabled, interval 300 msec, multiplier 5
  Transmit Delay is 1 sec, State BDR, Priority 1
  Designated Router (ID) 10.0.0.207, local address fe80::204:c0ff:fe22:73fe
  Backup Designated router (ID) 10.0.0.206, local address fe80::203:a0ff:fe9d:f3fe
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
    Hello due in 00:00:06
  Index 0/2/1, flood queue length 0
  Next 0(0)/0(0)/0(0)
  Last flood scan length is 2, maximum is 9
  Last flood scan time is 0 msec, maximum is 1 msec
  Neighbor Count is 1, Adjacent neighbor count is 1
    Adjacent with neighbor 10.0.0.207  (Designated Router)
  Suppress hello for 0 neighbor(s)
```

This table describes the significant fields shown in the display.

Table 130: show ospfv3 interface Field Descriptions

Field	Description
GigabitEthernet	Status of the physical link and operational status of the protocol.
Link Local Address	Interface link local address and interface ID.
Area	OSPFv3 area ID, process ID, instance ID, and router ID.
Transmit Delay	Transmit delay and interface state.
Designated Router	Designated router ID and respective interface IPv6 address.
Backup Designated router	Backup designated router ID and respective interface IPv6 address.
Timer intervals configured	Configuration of timer intervals.
Hello	Number of seconds until next hello packet is sent over this interface.
Index 0/2/1	Link, area and autonomous system flood indexes, and number of flood queue entries.
Next 0(0)/0(0)/0(0)	Next link, area and autonomous system flood information, data pointer, and index.
Last flood scan length	Length of last flood scan.
Last flood scan time	Time of last flood scan (in milliseconds).

Field	Description
Neighbor Count	Count of network neighbors and list of adjacent neighbors.
Suppress hello	Count of neighbors suppressing hello messages.

Related Commands

Command	Description
router ospfv3, on page 1367	Configures an OSPFv3 routing process.

show ospfv3 message-queue

To display the information about the queue dispatch values, peak lengths, and limits, use the **show ospfv3 message-queue** command in EXEC mode.

show ospfv3 [*process-name*] **message-queue**

This command has no keywords or arguments.

Command Default

None

Command Modes

EXEC

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 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
ospf	read

This is sample output from the **show ospfv3 message-queue** command:

```
RP/0/RSP0/CPU0:router#show ospfv3 message-queue
Mon May 31 16:07:47.143 CEST

OSPFv3 Process 0
  Hello Thread Packet Input Queue:
    Current queue length:         0
    Peak queue length:            2
    Queue limit:                  5000
    Packets received:             104091
    Packets processed:            104091
    Packets dropped:              0
    Processing quantum:           10
    Full quantum used:            0
    Pulses sent:                  104089
    Pulses received:              104089

  Router Thread Message Queue
    Current queue length:         0
    Peak queue length:            2
    Low queue limit:              8000
    Medium queuing limit:         9000
    High queuing limit:           9500
    Messages queued:              1472
    Messages deleted:             0
    Messages processed:           1472
```

```
Low queue drops:              0
Medium queue drops:           0
High queue drops:             0
Processing quantum:           300
Full quantum used:            0
Pulses sent:                  1484
Pulses received:              1484
```

show ospfv3 neighbor

To display Open Shortest Path First Version 3 (OSPFv3) neighbor information on an individual interface basis, use the **show ospfv3 neighbor** command in EXEC mode.

show ospfv3 [*process-name*] [*area-id*] **neighbor** [*type interface-path-id*] [*neighbor-id*] [**detail**]

Syntax Description

process-name	(Optional) Name that uniquely identifies an OSPFv3 routing process. The process name is defined by the **router ospfv3** command. If this argument is included, only information for the specified routing process is displayed.
area-id	(Optional) Area ID. If you do not specify an area, all areas are displayed.
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.
neighbor-id	(Optional) Neighbor router ID.
detail	(Optional) Displays all neighbors given in detail (lists all neighbors).

Command Default

No default behavior or values

Command Modes

EXEC

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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **show ospfv3 neighbor** command when the adjacency between two neighboring routers is not forming. Adjacency means that the routers synchronize their databases when they discover each other.

Task ID

Task ID	Operations
ospf	read

Examples

The following is sample output from the **show ospfv3 neighbor** command showing two lines of summary information for each neighbor:

```
RP/0/RSP0/CPU0:router# show ospfv3 neighbor

  Neighbors for OSPFv3 1

  Neighbor ID     Pri   State           Dead Time   Interface ID    Interface
  10.0.0.207      1     FULL/  -        00:00:35    3               GigabitEthernet 0/3/0/0

      Neighbor is up for 01:08:05
  10.0.0.207      1     FULL/DR         00:00:35    2               Ethernet0/0/0/0
      Neighbor is up for 01:08:05

  Total neighbor count: 2
```

This table describes the significant fields shown in the display.

Table 131: show ospfv3 neighbor Field Descriptions

Field	Description
ID	Neighbor router ID.
Pri	Router priority for designated router election. A router with a priority of 0 is never elected as the designated router or backup designated router.
State	OSPFv3 state.
Dead Time	Time (in hh:mm:ss) to elapse before OSPFv3 declares the neighbor dead.
Interface ID	Number that uniquely identifies an interface on a router.
Interface	Name of the interface that connects to this neighbor.
Neighbor is up	Time (in hh:mm:ss) that the OSPFv3 neighbor has been up.

The following is sample output showing summary information about the neighbor that matches the neighbor ID:

```
RP/0/RSP0/CPU0:router# show ospfv3 neighbor 10.0.0.207

  Neighbors for OSPFv3 1

   Neighbor 10.0.0.207
      In the area 0 via interface Ethernet0/0/0/0
      Neighbor: interface-id 2, link-local address fe80::204:c0ff:fe22:73fe
      Neighbor priority is 1, State is FULL, 6 state changes
      DR is 10.0.0.207 BDR is 10.0.0.206
      Options is 0x13
      Dead timer due in 00:00:38
      Neighbor is up for 01:09:21
      Index 0/1/2, retransmission queue length 0, number of retransmission 1
      First 0(0)/0(0)/0(0) Next 0(0)/0(0)/0(0)
      Last retransmission scan length is 1, maximum is 1
```

```
    Last retransmission scan time is 0 msec, maximum is 0 msec

 Neighbor 10.0.0.207
    In the area 1 via interface GigabitEthernet 0/3/0/0
    Neighbor: interface-id 3, link-local address fe80::3034:30ff:fe33:3742
    Neighbor priority is 1, State is FULL, 6 state changes
    Options is 0x13
    Dead timer due in 00:00:38
    Neighbor is up for 01:09:21
    Index 0/1/1, retransmission queue length 0, number of retransmission 1
    First 0(0)/0(0)/0(0) Next 0(0)/0(0)/0(0)
    Last retransmission scan length is 1, maximum is 1
    Last retransmission scan time is 0 msec, maximum is 0 msec


Total neighbor count: 2
```

This table describes the significant fields shown in the display.

Table 132: show ospfv3 neighbor 10.0.0.207 Field Descriptions

Field	Description
Neighbor	Neighbor router ID.
In the area	Area and interface through which the OSPFv3 neighbor is known.
link-local address	Link local address of the interface.
Neighbor priority	Router priority of neighbor and neighbor state.
State	OSPFv3 state.
state changes	Number of state changes for this neighbor.
DR is	Neighbor ID of the designated router.
BDR is	Neighbor ID of the backup designated router.
Options	Hello packet options field contents (E-bit only; possible values are 0 and 2; 2 indicates that area is not a stub; 0 indicates that area is a stub).
Dead timer	Time (in hh:mm:ss) to elapse before OSPFv3 declares the neighbor dead.
Neighbor is up	Time (in hh:mm:ss) that OSPFv3 neighbor has been up.
Index	Index and the remaining lines of this command give detailed information about flooding information received from the neighbor.

The following sample output shows the neighbors that match the neighbor ID on the interface when the interface along with the neighbor ID is specified:

```
RP/0/RSP0/CPU0:router# show ospfv3 neighbor GigabitEthernet 0/3/0/1 10.0.0.207

  Neighbors for OSPFv3 1
```

```
 Neighbor 10.0.0.207
    In the area 0 via interface GigabitEthernet 0/3/0/1
    Neighbor: interface-id 2, link-local address fe80::204:c0ff:fe22:73fe
    Neighbor priority is 1, State is FULL, 6 state changes
    DR is 10.0.0.207 BDR is 10.0.0.206
    Options is 0x13
    Dead timer due in 00:00:39
    Neighbor is up for 01:11:21
    Index 0/1/2, retransmission queue length 0, number of retransmission 1
    First 0(0)/0(0)/0(0) Next 0(0)/0(0)/0(0)
    Last retransmission scan length is 1, maximum is 1
    Last retransmission scan time is 0 msec, maximum is 0 msec


Total neighbor count: 1
```

This table describes the significant fields shown in the display.

Table 133: show ospfv3 neighbor GigabitEthernet 0/3/0/1 10.0.0.207 Field Descriptions

Field	Description
Neighbor	Neighbor router ID.
In the area	Area and interface through which the OSPFv3 neighbor is known.
link-local address	Link local address of the interface.
Neighbor priority	Router priority of neighbor and neighbor state.
State	OSPFv3 state.
state changes	Number of state changes for this neighbor.
DR is	Neighbor ID of the designated router.
BDR is	Neighbor ID of the backup designated router.
Options	Hello packet options field contents (E-bit only; possible values are 0 and 2; 2 indicates that area is not a stub; 0 indicates that area is a stub).
Dead timer	Time (in hh:mm:ss) to elapse before OSPFv3 declares the neighbor dead.
Neighbor is up	Time (in hh:mm:ss) that OSPFv3 neighbor has been up.
Index	Index and the remaining lines of this command give detailed information about flooding information received from the neighbor.

The following sample output shows all neighbors on the interface when the interface is specified:

```
RP/0/RSP0/CPU0:router# show ospfv3 neighbor GigabitEthernet 0/3/0/1

  Neighbors for OSPFv3 1

  Neighbor ID     Pri   State           Dead Time   Interface ID    Interface
  10.0.0.207      1     FULL/DR         00:00:37    2               GigabitEthernet 0/3/0/1
```

```
    Neighbor is up for 01:12:33

Total neighbor count: 1
```

This table describes the significant fields shown in the display.

Table 134: show ospfv3 neighbor GigabitEthernet 0/3/0/1 Field Descriptions

Field	Description
Neighbor ID	Neighbor router ID.
Pri	Router priority for designated router election. A router with a priority of 0 is never elected as the designated router or backup designated router.
State	OSPF state.
Dead Time	Time (in hh:mm:ss) to elapse before OSPF declares the neighbor dead.
Interface ID	Number that uniquely identifies an interface on a router.
Interface	Name of the interface that connects to this neighbor.
Neighbor is up	Amount of time (in hh:mm:ss) that the OSPF neighbor has been up.

The following is sample output showing detailed neighbor information for GigabitEthernet interface 0/3/0/1:

```
RP/0/RSP0/CPU0:router# show ospfv3 neighbor GigabitEthernet 0/3/0/1 detail

  Neighbors for OSPFv3 1

   Neighbor 10.0.0.207
      In the area 0 via interface GigabitEthernet 0/3/0/1
      Neighbor: interface-id 2, link-local address fe80::204:c0ff:fe22:73fe
      Neighbor priority is 1, State is FULL, 6 state changes
      DR is 10.0.0.207 BDR is 10.0.0.206
      Options is 0x13
      Dead timer due in 00:00:39
      Neighbor is up for 01:13:40
      Index 0/1/2, retransmission queue length 0, number of retransmission 1
      First 0(0)/0(0)/0(0) Next 0(0)/0(0)/0(0)
      Last retransmission scan length is 1, maximum is 1
      Last retransmission scan time is 0 msec, maximum is 0 msec

  Total neighbor count: 1
```

This table describes the significant fields shown in the display.

Table 135: show ospfv3 neighbor GigabitEthernet 0/3/0/1 detail Field Descriptions

Field	Description
Neighbor	Neighbor router ID.

Field	Description
In the area	Area and interface through which the OSPFv3 neighbor is known.
link-local address	Link local address of the interface.
Neighbor priority	Router priority of neighbor and neighbor state.
State	OSPFv3 state.
state changes	Number of state changes for this neighbor.
DR is	Neighbor ID of the designated router.
BDR is	Neighbor ID of the backup designated router.
Options	Hello packet options field contents (E-bit only; possible values are 0 and 2; 2 indicates that area is not a stub; 0 indicates that area is a stub).
Dead timer	Time (in hh:mm:ss) to elapse before OSPFv3 declares the neighbor dead.
Neighbor is up	Time (in hh:mm:ss) that the OSPFv3 neighbor has been up.
Index	Index and the remaining lines of this command give detailed information about flooding information received from the neighbor.

Related Commands

Command	Description
router ospfv3, on page 1367	Configures an OSPFv3 routing process.

show ospfv3 request-list

To display the first ten link-state requests pending that the local router is making to the specified Open Shortest Path First Version 3 (OSPFv3) neighbor and interface, use the **show ospfv3 request-list** command in EXEC mode.

show ospfv3 [*process-name*] [*area-id*] **request-list** [*type interface-path-id*] [*neighbor-id*]

Syntax Description

process-name	(Optional) Name that uniquely identifies an OSPFv3 routing process. The process name is defined by the **router ospfv3** command. If this argument is included, only information for the specified routing process is displayed.
area-id	(Optional) Area ID. If you do not specify an area, all areas are displayed.
type	(Optional) Interface type. For more information, use the question mark (?) online help function.
interface-path-id	(Optional) 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.
neighbor-id	(Optional) Router ID of the OSPFv3 neighbor. This argument must be in 32-bit dotted-decimal notation, similar to an IPv4 address.

Command Default

No default behavior or values

Command Modes

EXEC

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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

You might use this command when the databases of two neighboring routers are out of synchronization or if the adjacency is not forming between them. Adjacency means that the routers synchronize their databases when they discover each other.

You can look at the list to determine if one router is trying to request a particular database update. Entries that appear to be suspended in the list usually indicate that updates are not being delivered. One possible reason for this behavior is a maximum transmission unit (MTU) mismatch between the routers.

You might also look at this list to make sure it is not corrupted. The list should refer to database entries that actually exist.

Task ID

Task ID	Operations
ospf	read

Examples

The following sample output shows request lists for neighbor 10.0.0.207 on the OSPFv3 1 process:

```
RP/0/RSP0/CPU0:router# show ospfv3 1 request-list 10.0.0.207 GigabitEthernet 0/3/0/0

  Request Lists for OSPFv3 1

   Neighbor 10.0.0.207, interface GigabitEthernet 0/3/0/0 address fe80::3034:30ff:fe33:3742


    Type   LS ID            ADV RTR         Seq NO         Age   Checksum
       1   192.168.58.17     192.168.58.17   0x80000012     12   0x0036f3
       2   192.168.58.68     192.168.58.17   0x80000012     12   0x00083f
```

This table describes the significant fields shown in the display.

Table 136: show ospfv3 request-list Field Descriptions

Field	Description
Neighbor	Router ID of the neighboring router.
interface	Name of the interface that connects to this neighbor.
address	IPv6 address of the neighbor.
Type	Type of link-state advertisement (LSA).
LS ID	Link-state ID of the LSA.
ADV RTR	Router ID of the advertising router.
Seq NO	Sequence number of the LSA.
Age	Age of the LSA (in seconds).
Checksum	Checksum of the LSA.

Related Commands

Command	Description
router ospfv3, on page 1367	Configures an OSPFv3 routing process.
show ospfv3 retransmission-list, on page 1402	Displays the first ten link-state entries in the retransmission list that the local router sends to the specified neighbor over the specified interface.

show ospfv3 retransmission-list

To display the first ten link-state entries in the retransmission list that the local router sends to the specified neighbor over the specified interface, use the **show ospfv3 retransmission-list** command in EXEC mode.

show ospfv3 [*process-name*] [*area-id*] **retransmission-list** [*type interface-path-id*] [*neighbor-id*]

Syntax Description

process-name	(Optional) Name that uniquely identifies an Open Shortest Path First Version 3 (OSPFv3) routing process. The process name is defined by the **router ospfv3** command. If this argument is included, only information for the specified routing process is displayed.
area-id	(Optional) Area ID. If you do not specify an area, all areas are displayed.
type	(Optional) Interface type. For more information, use the question mark (?) online help function.
interface-path-id	(Optional) 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.
neighbor-id	(Optional) IP address of the OSPFv3 neighbor.

Command Default

No default behavior or values

Command Modes

EXEC

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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

You might use this command when the databases of two neighboring routers are out of synchronization or if the adjacency is not forming between them. Adjacency means that the routers synchronize their databases when they discover each other.

You can look at the list to determine if one router is trying to request a particular database update. Entries that appear to be suspended in the list usually indicate that updates are not being delivered. One possible reason for this behavior is a maximum transmission unit (MTU) mismatch between the routers.

You might also look at this list to make sure it is not corrupted. The list should refer to database entries that actually exist.

Task ID

Task ID	Operations
ospf	read

Examples

The following sample output shows the retransmission list for neighbor 10.0.124.4 on GigabitEthernet interface 0/3/0/0:

```
RP/0/RSP0/CPU0:router#show ospfv3 retransmission-list 10.0.124.4 GigabitEthernet 0/3/0/0

  Neighbor 10.0.124.4, interface GigabitEthernet 0/3/0/0 address fe80::3034:30ff:fe33:3742
```

This table describes the significant fields shown in the display.

Table 137: show ospfv3 retransmission-list 10.0.124.4 GigabitEthernet 0/3/0/0 Field Descriptions

Field	Description
Neighbor	Router ID of the neighboring router.
interface	Name of the interface that connects to this neighbor.
address	IPv6 address of the neighbor.

Related Commands

Command	Description
router ospfv3, on page 1367	Configures an OSPFv3 routing process.
show ospfv3 request-list, on page 1400	Displays the first ten link-state requests pending that the local router is making to the specified neighbor and interface.

show ospfv3 routes

To display the Open Shortest Path First Version 3 (OSPFv3) route table, use the **show ospfv3 routes** command in EXEC mode.

show ospfv3 [*process-name*] **routes** [**external** | **connected**] [*ipv6-prefix*/*prefix-length*]
show ospfv3 [*process-name*] **routes summary**

Syntax Description

process-name	(Optional) Name that uniquely identifies an OSPFv3 routing process. The process name is defined by the **router ospf** command. If this argument is included, only information for the specified routing process is displayed.
external	(Optional) Displays routes redistributed from other protocols.
connected	(Optional) Displays connected routes.
ipv6-prefix	(Optional) IP Version 6 (IPv6) prefix, which limits output to a specific route. This argument must be in the form documented in RFC 2373, in which the address is specified in hexadecimal using 16-bit values between colons.
/ *prefix-length*	(Optional) Length of the IPv6 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.
summary	Displays a summary of the route table.

Command Default

No default behavior or values

Command Modes

EXEC

Command History

Release	Modification
Release 3.7.2	This command was introduced.
Release 3.9.0	Asplain format for 4-byte Autonomous system numbers notation was supported. The input parameters and output were modified to display 4-byte autonomous system numbers and extended communities in either asplain or asdot notations.

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.

Use the **show ospfv3 routes** command to display the OSPFv3 private routing table (which contains only routes calculated by OSPFv3). If there is something wrong with a route in the Routing Information Base (RIB), check the OSPFv3 copy of the route to determine if it matches the RIB contents. If it does not match, there is a synchronization problem between OSPFv3 and the RIB. If the routes match and the route is incorrect, OSPFv3 has made an error in its routing calculation.

Task ID

Task ID	Operations
ospf	read

Examples

The following sample output shows the route table for OSPFv3 process 1:

```
RP/0/RSP0/CPU0:router# show ospfv3 1 routes

  Route Table for OSPFv3 1 with ID 10.3.4.2

  * 3000:11:22::/64, Inter, cost 21/0, area 1
      GigabitEthernet 0/3/0/0, fe80::3034:30ff:fe33:3742
      10.0.0.207/200
  * 3000:11:22:1::/64, Inter, cost 31/0, area 1
  GigabitEthernet 0/3/0/0, fe80::3034:30ff:fe33:3742
      10.0.0.207/1
  * 3333::/56, Ext2, cost 20/1, P:0 F:0
  GigabitEthernet 0/3/0/0, fe80::3034:30ff:fe33:3742
      10.0.0.207/0
  * 6050::/56, Ext2, cost 20/1, P:0 F:0
  GigabitEthernet 0/3/0/0, fe80::3034:30ff:fe33:3742
      10.0.0.207/1
  * 7002::/56, Intra, cost 10/0, area 0
      Ethernet0/0/0/0, connected

  * 3000:11:22::/64, Inter, cost 21/0, area 1
  GigabitEthernet 0/3/0/0, fe80::3034:30ff:fe33:3742
      10.0.0.207/200
```

This table describes the significant fields shown in the display.

Table 138: show ospfv3 1 route Field Descriptions

Field	Description
3000:11:22::/64	Route prefix to the local router.
Inter	Prefix 3000:11:22::/64 is interarea.
cost 21/0	Sum of the link costs required to reach prefix 3000:11:22::/64. 0. In this example, 20 is the external cost.
GigabitEthernet 0/3/0/0	Packets destined for prefix 3000:11:22::/64 are sent over the GigabitEthernet 0/3/0/0 interface.
fe80::3034:30ff:fe33:3742	Next-hop router on the path to prefix 3000:11:22::/64.
10.0.0.207	Router 10.0.0.207 is the router that advertised this route.

Related Commands

Command	Description
router ospfv3, on page 1367	Configures an OSPFv3 routing process.

show ospfv3 statistics rib-thread

To display RIB thread statistics, use the **show ospfv3 statistics rib-thread** command in EXEC mode.

show ospfv3 [*process-name* [*area-id*]] **statistics rib-thread**

Syntax Description

process-name	(Optional) Name that uniquely identifies an OSPF routing process. The process name is defined by the **router ospfv3** command. If this argument is included, only information for the specified routing process is displayed.
area id	(Optional) Area number used to define the particular area.

Command Default

None

Command Modes

EXEC

Command History

Release	Modification
Release 4.1.0	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
ospf	read

This is sample output from **show ospfv3 statistics rib-thread** command:

```
Router#show ospfv3 0 statistics rib-thread
Mon May 10 17:48:29.011 CEST
OSPFv3 0 RIB thread statistics
 Queue statistics:
  Last entry dequeue           10127056 msecs ago (14:59:42.171)
  RIB thread active            NO
  Total RIB thread signals     30
  Current queue length         0
  Maximum queue length         2
  Total entries queued         31
  Total entries dequeued       31
  Maximum latency (msec)       5.000
  Average latency (msec)       0.323
 Queue errors:
  Enqueue errors               0
  Dequeue errors               0
 RIB batch statistics:
  Batches sent to RIB          31
  Batch all routes OK          31
  Batch some routes backup     0
```

```
RIB batch errors:
 Batches version mismatch     0
 Batches missing connection   0
 Batches no table             0
 Batch route table limit      0
 Batch route errors           0
 Batch errors                 0
 Route table limit            0
 Route path errors            0
 Route errors                 0
 Path table limit             0
 Path errors                  0
```

show ospfv3 summary-prefix

To display Open Shortest Path First Version 3 (OSPFv3) aggregated summary address information, use the **show ospfv3 summary-prefix** command in EXEC mode.

show ospfv3 [*process-name*] **summary-prefix**

Syntax Description

process-name	(Optional) Name that uniquely identifies an OSPFv3 routing process. The process name is defined by the **router ospfv3** command. If this argument is included, only information for the specified routing process is displayed.

Command Default

No default behavior or values

Command Modes

EXEC

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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **show ospfv3 summary-prefix** command if you configured summarization of external routes with the **summary-prefix** command and you want to display configured summary addresses.

Task ID

Task ID	Operations
ospf	read

Examples

The following sample output shows the summary prefix address for the OSPFv3 1 process:

```
RP/0/RSP0/CPU0:router# show ospfv3 1 summary-prefix

  OSPFv3 Process 1, Summary-prefix

  4004:f000::/32 Metric 20, Type 2, Tag 0
```

This table describes the significant fields shown in the display.

Table 139: show ospfv3 1 summary-prefix Field Descriptions

Field	Description
4004:f000::/32	Summary prefix designated for a range of IPv6 prefixes. The length of the IPv6 prefix.
Metric	Metric used to advertise the summary routes.

Field	Description
Type	External link-state advertisements (LSAs) metric type.
Tag	Tag value that can be used as a “match” value for controlling redistribution through route maps.

Related Commands

Command	Description
router ospfv3, on page 1367	Configures an OSPFv3 routing process.
summary-prefix (OSPFv3), on page 1424	Creates aggregate addresses for routes being redistributed from another routing protocol into OSPFv3.

show ospfv3 virtual-links

To display parameters and the current state of Open Shortest Path First Version 3 (OSPFv3) virtual links, use the **show ospfv3 virtual-links** command in EXEC mode.

show ospfv3 [*process-name*] **virtual-links**

Syntax Description

process-name	(Optional) Name that uniquely identifies an OSPFv3 routing process. The process name is defined by the **router ospfv3** command. If this argument is included, only information for the specified routing process is displayed.

Command Default

No default behavior or values

Command Modes

EXEC

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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

The information displayed in the **show ospfv3 virtual-links** command is useful in debugging OSPFv3 routing operations.

Task ID

Task ID	Operations
ospf	read

Examples

The following sample output shows the virtual links for the OSPFv3 1 process:

```
RP/0/RSP0/CPU0:router# show ospfv3 1 virtual-links

  Virtual Links for OSPFv3 1
  Virtual Link to router 172.31.101.2 is up
    Interface ID 16, IPv6 address 3002::206
    Transit area 0.0.0.1, via interface GigabitEthernet 0/3/0/0, Cost of using 11
    Transmit Delay is 5 sec, State POINT_TO_POINT,
    Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
    Hello due in 0:00:08
    Adjacency State FULL
```

This table describes the significant fields shown in the display.

Table 140: show ospfv3 virtual-links Field Descriptions

Field	Description
Virtual Link to router is up	Specifies the OSPFv3 neighbor, and if the link to that neighbor is up or down.
Interface ID	ID of the virtual link interface.
IPv6 address	IPv6 address of virtual link endpoint.
Transit area	Transit area through which the virtual link is formed.
via interface	Interface through which the virtual link is formed.
Cost	Cost of reaching the OSPF neighbor through the virtual link.
Transmit Delay	Transmit delay on the virtual link.
State POINT_TO_POINT	State of the OSPFv3 neighbor.
Timer intervals	Various timer intervals configured for the link.
Hello due in	When the next hello message is expected from the neighbor (in hh:mm:ss).
Adjacency State	Adjacency state between the neighbors.

Related Commands

Command	Description
router ospfv3, on page 1367	Configures an OSPFv3 routing process.

show protocols (OSPFv3)

To display information about the Open Shortest Path First Version 3 (OSPFv3) process running on the router, use the **show protocols** command in EXEC mode.

show protocols [**afi-all** | **ipv4** | **ipv6**] [**all***protocol*]

Syntax Description

afi-all	(Optional) Specifies all address families.
ipv4	(Optional) Specifies an IPv4 address family.
ipv6	(Optional) Specifies an IPv6 address family.
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** • **eigrp** • **isis** • **ospf** • **rip** For the IPv6 address family, the options are: • **bgp** • **eigrp** • **isis** • **ospfv3**

Command Default

The default address family is IPv4.

Command Modes

EXEC

Command History

Release	Modification
Release 3.7.2	This command was introduced.
Release 3.9.0	Asplain format for 4-byte Autonomous system numbers notation was supported. The input parameters and output were modified to display 4-byte autonomous system numbers and extended communities in either asplain or asdot notations.

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
ospf	read

Examples

The following example is sample output from the **show protocols** command:

```
RP/0/RSP0/CPU0:router# show protocols ipv6 ospfv3

  Routing Protocol OSPFv3 1
    Router Id:10.0.0.1
    Distance:110
    Redistribution:
      None
    Area 0
      GigabitEthernet 0/2/0/2
      Loopback1
```

This table describes the significant fields shown in the display.

Table 141: show protocols Field Descriptions

Field	Description
Router Id	Router ID of the OSPFv3 process.
Distance	Administrative distance for the protocol. This distance determines the priority the Routing Information Base (RIB) gives to the routes, as opposed to other protocols, for example, IS-IS.
Redistribution	Protocols from which this OSPFv3 process is redistributing routes.
Area	OSPFv3 areas defined in this process, followed by their associated interfaces.

snmp context (OSPFv3)

To specify an SNMP context for an OSPFv3 instance, use the **snmp context** command in router configuration mode or in VRF configuration mode. To remove the SNMP context, use the **no** form of this command.

snmp context *context_name*
no snmp context *context_name*

Syntax Description

context_name	Specifies name of the SNMP context for OSPFv3 instance.

Command Default

SNMP context is not specified.

Command Modes

Router OSPFv3 configuration

VRF configuration

Command History

Release	Modification
Release 4.2.1	This command was introduced.

Usage Guidelines

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

The snmp-server commands need to be configured to perform SNMP request for the OSPF instance. Refer *SNMP Server Commands* module in *System Management Command Reference for Cisco ASR 9000 Series Routers* for information on using the snmp-server commands.

Note To map an SNMP context with a protocol instance, topology or VRF entity, use the **snmp-server context mapping** command. However, the **feature** option of this command does not work with OSPFv3 protocol.

Task ID

Task ID	Operation
ospf	read, write

This example shows how to configure an SNMP context *foo* for OSPFv3 instance *100*:

```
Router#configure
Router(config)#router ospfv3 100
Router(config-ospf)#snmp context foo
```

This example shows how to configure **snmp-server** commands to be used with the **snmp context** command:

```
Router(config)#snmp-server host 10.0.0.2 traps version 2c public udp-port 1620
Router(config)#snmp-server community public RW
Router(config)#snmp-server contact foo
Router(config)#snmp-server community-map public context foo
```

This is a sample SNMP context configuration for OSPFv3 instance *100*:

```
snmp-server host 10.0.0.2 traps version 2c public udp-port 1620
snmp-server community public RW
snmp-server contact foo

snmp-server community-map public context foo


router ospfv3 100
 router-id 2.2.2.2
 bfd fast-detect
 nsf cisco
snmp context foo
 area 0
  interface Loopback1
  !
 !
 area 1
  interface GigabitEthernet0/2/0/1
   demand-circuit enable
  !
  interface POS0/3/0/0
  !
  interface POS0/3/0/1
  !
 !
!
```

Related Commands

Command	Description
Test	Enables SNMP trap for an OSPFv3 instance.
snmp-server host	Specifies the recipient of an SNMP notification operation.
snmp-server community	Configures the community access string to permit access to the Simple Network Management Protocol (SNMP).
snmp-server contact	Sets the Simple Network Management Protocol (SNMP) system contact.
snmp-server community-map	Associates a Simple Network Management Protocol (SNMP) community with an SNMP context.

snmp trap (OSPFv3)

To enable SNMP trap for an OSPFv3 instance, use the **snmp trap** command in VRF configuration mode. To disable SNMP trap for the OSPFv3 instance, use the **no** form of this command.

snmp trap
no snmp trap

Syntax Description

This command has no keywords or arguments.

Command Default

Disabled.

Command Modes

OSPFv3 VRF configuration

Command History

Release	Modification
Release 4.2.1	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
ospf	read, write

This example shows how to enable SNMP trap for OSPFv3 instance *100* under VRF *vrf-1*:

```
Router#configure
Router(config)#router ospfv3 100
Router(config-ospf)#vrf vrf-1
Router(config-ospf-vrf)#snmp trap
```

Related Commands

Command	Description
snmp context ospfv3	

snmp trap rate-limit (OSPFv3)

To control the number of traps that OSPFv3 sends by configuring window size and the maximum number of traps during that window, use the **snmp trap rate-limit** command in router OSPFv3 configuration mode or OSPFv3 VRF configuration mode. To disable configuring the window size and maximum number of traps during the window, use the **no** form of this command.

snmp trap rate-limit *window-size max-num-traps*
no snmp trap rate-limit *window-size max-num-traps*

Syntax Description

window-size	Specifies the trap rate limit sliding window size. The range is 2 to 60 windows.
max-num-traps	Specifies the maximum number of traps sent in window time. The range is 0 to 300 traps.

Command Default

None

Command Modes

Router OSPFv3 configuration

OSPFv3 VRF configuration

Command History

Release	Modification
Release 4.2.1	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
ospf	read, write

This example shows how to set the trap rate limit sliding window size to 50 and the maximum number of traps sent to 250 for OSPFv3 instance *100* under vrf *vrf1*:

```
Router#configure
Router(config)#router ospfv3 100
Router(config-ospfv3)#vrf vrf1
Router(config-ospfv3-vrf)#snmp trap rate-limit 50 250
```

spf prefix-priority (OSPFv3)

To prioritize OSPFv3 prefix installation into the global Routing Information Base (RIB) during Shortest Path First (SPF) run, use the **spf prefix-priority** command in router configuration mode or VRF configuration mode. To return to the system default value, use the **no** form of this command.

spf prefix-priority route-policy *policy-name* [**disable**]
spf prefix-priority route-policy *policy-name*

Syntax Description

route-policy	Specifies the route-policy to prioritize route installation.
policy-name	Name of the route policy.
disable	Disables SPF prefix priority

Command Default

SPF prefix prioritization is disabled.

Command Modes

Router configuration

VRF configuration

Command History

Release	Modification
Release 4.2.1	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
ospf	read, write

This example shows how to configure OSPFv3 SPF prefix prioritization:

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# prefix-set ospf3-critical-prefixes
RP/0/RP0/CPU0:router(config-pfx)# 66.0.0.0/16
RP/0/RP0/CPU0:router(config-pfx)# end-set
RP/0/RP0/CPU0:router(config)# route-policy ospf3-spf-priority
RP/0/RP0/CPU0:router(config-rpl)# if destination in ospf-critical-prefixes then set
spf-priority critical
endif
RP/0/RP0/CPU0:router(config-rpl)# end-policy
RP/0/RP0/CPU0:router(config-rpl)# commit
RP/0/RP0/CPU0:router(config-rpl)# exit
RP/0/RP0/CPU0:router(config)# router ospfv3 1
RP/0/RP0/CPU0:router(config-ospf)# router-id 66.0.0.1
```

```
RP/0/RP0/CPU0:router(config-ospf)# spf prefix-priority route-policy ospf-spf-priority
```

Related Commands

Command	Description
prefix-set	Enters prefix set configuration mode and defines a prefix set.
route-policy (RPL)	Defines a route policy and enters route-policy configuration mode.

stub (OSPFv3)

To define an area as a stub area for Open Shortest Path First Version 3 (OSPFv3), use the **stub** command in area configuration mode. To disable this function, use the **no** form of this command.

stub [**no-summary**]
no stub

Syntax Description

no-summary	(Optional) Prevents an area border router (ABR) from sending summary link advertisements into the stub area. Areas with this option are known as *totally stubby* areas.

Command Default

No stub area is defined.

Command Modes

Area 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

You must configure the **stub** command on all routers in the stub area. Use the **default-cost area** command on the ABR of a stub area to specify the cost of the default route advertised into the stub area by the ABR.

Two stub area router configuration commands exist: the **stub** and **default-cost** commands. In all routers attached to the stub area, the area should be configured as a stub area using the **stub** command. Use the **default-cost** command only on an ABR attached to the stub area. The **default-cost** command provides the metric for the summary default route generated by the ABR into the stub area.

To further reduce the number of link-state advertisements (LSAs) sent into a stub area, you can configure the **no-summary** keyword on the ABR to prevent it from sending summary LSAs (LSA Type 3) into the stub area.

A stub area does not accept information about routes external to the autonomous system.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to create stub area 5 and specifies a cost of 20 for the default summary route sent into this stub area:

```
RP/0/RSP0/CPU0:router(config)# router ospfv3 201
RP/0/RSP0/CPU0:router(config-ospfv3)# area 5
RP/0/RSP0/CPU0:router(config-ospfv3-ar)# stub
```

```
RP/0/RSP0/CPU0:router(config-ospfv3-ar)# default-cost 20
```

Related Commands

Command	Description
default-cost (OSPFv3), on page 1308	Specifies a cost for the default summary route sent into a stub area.

stub-router

To modify self originated router LSAs when stub router is active, use the **stub-router** command in an appropriate configuration mode. To disable this function, use the **no** form of this command.

stub-router router-lsa [**r-bit** | **v6-bit** | **max-metric**] [**always**] [**on-proc-migration** *interval*] [**on-proc-restart** *interval*] [**on-switchover** *interval*] [**on-startup** [*interval* | **wait-for-bgp**]] [**summary-lsa** [*metric*]] [**external-lsa** [*metric*]] [**include-stub**]
stub-router router-lsa [**r-bit** | **v6-bit** | **max-metric**]

Syntax Description

router-lsa	Specifies that always originate router link-state advertisements (LSAs) with the stub-router.
r-bit	Router-LSAs are originated with R-bit clear (v6 bit set), which means the node does not act as a transit router. Directly connected networks (native to OSPF) are still reachable within the OSPF area.
v6-bit	Router-LSAs are originated with V6 bit clear (and also r-bit clear). That means the node is not willing to receive any ipv6 traffic. Other ospfv3 routers won't install any route to a node with v6-bit clear.
max-metric	Router-LSAs are originated with maximum metric. Unlike the r-bit and v6-bit mode, the router may still act as a transit node, if there is no alternate path.
always	Stub-router mode is activated unconditionally.
on-proc-migration	Stub-router mode is activated for the desired period of time, upon ospfv3 process migration.
on-proc-restart	Stub-router mode is activated for the desired period of time, upon ospfv3 process restart.
on-switchover	Stub-router mode is activated for the desired period of time, upon RP failover.
on-startup	Stub-router mode is activated (for configured time, or until BGP converges) upon router startup (boot).
wait-for-bgp	Stub-router mode is terminated upon BGP convergence in ipv6 unicast address family. This option could only be used in the global routing table, not in a non-default VRF. This option is only supported with the on-startup trigger when the router boots.
summary-lsa	If enabled, summary LSAs are advertised with modified metric when stub-router is active. This configuration is applicable to max-metric mode. In r-bit mode, ABR/ASBR functionality is implicitly disabled and routers will not use this node as an ABR/ASBR, since it declares no transit capability (r-bit clear). If enabled and metric is not explicitly configured, the default metric for summary LSAs when stub-router active is 16711680 (0xFF0000).

external-lsa	If enabled, external LSAs are advertised with modified metric when stub-router is active. This configuration is applicable to max-metric mode. In r-bit mode, ABR/ASBR functionality is implicitly disabled and routers will not use this node as an ABR/ASBR, since it declares no transit capability (r-bit clear). If enabled and metric is not explicitly configured, the default metric for external LSAs when stub-router active is 16711680 (0xFF0000).
include-stub	If enabled, intra-area-prefix LSAs that are referencing router LSA are advertised with maximum metric (0xffff) when stub-router is active. Intra-area-prefix LSAs that are referencing network LSA do not change metric Can be used in r-bit and max-metric modes. /128 prefixes that are normally advertised with LA-bit set and 0 metric are also advertised with maximum metric and LA-bit clear when stub-router is active.

Command Default

Disabled.

Command Modes

Router OSPFv3 configuration

OSPFv3 VRF configuration

Command History

Release	Modification
Release 4.2.0	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.

Only one method (r-bit, v6-bit, max-metric) could be activated at a time. Configuring the methods simultaneously, or different method per trigger, is not supported.

Task ID

Task ID	Operation
ospf	read, write

This example shows how to configure router LSAs are originated with R-bit clear under OSPFv3 VRF, *vrf_1*:

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#router ospfv3 1
RP/0/RSP0/CPU0:router(config-ospfv3)#vrf vrf_1
RP/0/RSP0/CPU0:router(config-ospfv3-vrf)#stub-router router-lsa r-bit
```

summary-prefix (OSPFv3)

To create aggregate addresses for routes being redistributed from another routing protocol into Open Shortest Path First Version 3 (OSPFv3) protocol, use the **summary-prefix** command in an appropriate configuration mode. To stop summarizing redistributed routes, use the **no** form of the command.

summary-prefix *ipv6-prefix/prefix-length* [**not-advertise**] **tag** *tag*
no summary-prefix *ipv6-prefix/prefix-length*

Syntax Description

ipv6-prefix	Summary prefix designated for a range of IP Version 6 (IPv6) prefixes. This argument must be in the form documented in RFC 2373, where the address is specified in hexadecimal using 16-bit values between colons.
/ *prefix-length*	Length of the IPv6 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.
not-advertise	(Optional) Suppresses summary routes that match the address and mask pair from being advertised.
tag *tag*	(Optional) Specifies a tag value that can be used as a "match" value for controlling redistribution.

Command Default

When this command is not used in router configuration mode, aggregate addresses are not created for routes being redistributed from another routing protocol into the OSFPv3 protocol.

Command Modes

Router OSPFv3 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **summary-prefix** command to cause an OSPFv3 Autonomous System Boundary Router (ASBR) to advertise one external route as an aggregate for all redistributed routes that are covered by the address. This command summarizes only routes from other routing protocols that are being redistributed into OSPFv3.

You can use this command multiple times to summarize multiple groups of addresses. The metric used to advertise the summary is the lowest metric of all the more specific routes. This command helps reduce the size of the routing table.

If you want to summarize routes between OSPFv3 areas, use the **range** command.

Task ID

Task ID	Operations
ospf	read, write

Examples

In the following example, if summary prefix 4004:f000:132 is configured and routes 4004:f000:1::/64, 4004:f000:2::/64, and 4004:f000:3::/64 are redistributed into OSPFv3; only route 4004:f000::/32 is advertised in an external link-state advertisement:

```
RP/0/RSP0/CPU0:router(config-ospfv3)# summary-prefix 4004:f000::/32
```

Related Commands

Command	Description
range (OSPFv3), on page 1357	Consolidates and summarizes routes at an area boundary.

timers lsa arrival

To set the minimum interval at which the software accepts the same link-state advertisement (LSA) from Open Shortest Path First Version 3 (OSPFv3) neighbors, use the **timers lsa arrival** command in an appropriate configuration mode. To restore the default value, use the **no** form of this command.

timers lsa arrival *milliseconds*
no timers lsa arrival

Syntax Description

milliseconds	Minimum delay (in milliseconds) that must pass between acceptance of the same LSA arriving from neighbors. Range is 0 to 60000 milliseconds.

Command Default

1000 milliseconds

Command Modes

Router OSPFv3 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **timers lsa arrival** command to control the minimum interval for accepting the same LSA. The same LSA is an LSA instance that contains the same LSA ID number, LSA type, and advertising router ID. If an instance of the same LSA arrives sooner than the interval that is set, the LSA is dropped.

We recommended that the *milliseconds* value of the **timers lsa arrival** command be less than or equal to the *hold-interval* value of the **timers throttle lsa all** command for the neighbor.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to set the minimum interval for accepting the same LSA at 2000 milliseconds:

```
RP/0/RSP0/CPU0:router(config)# router ospfv3 1
RP/0/RSP0/CPU0:router(config-ospfv3)# timers throttle lsa all 200 10000 45000
RP/0/RSP0/CPU0:router(config-ospfv3)# timers lsa arrival 2000
```

Related Commands

Command	Description
timers throttle lsa all (OSPFv3), on page 1433	Sets rate-limiting values for LSAs being generated.

timers pacing flood

To configure link-state advertisement (LSA) flood packet pacing, use the **timers pacing flood** command in an appropriate configuration mode. To restore the default flood packet pacing value, use the **no** form of this command.

timers pacing flood *milliseconds*
no timers pacing flood

Syntax Description

milliseconds	Time (in milliseconds) at which LSAs in the flooding queue are paced in between updates. Range is 5 milliseconds to 100 milliseconds.

Command Default

milliseconds: 33

Command Modes

Router OSPFv3 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Configuring OSPFv3 flood pacing timers allows you to control interpacket spacing between consecutive link-state update packets in the OSPF transmission queue. Use the **timers pacing flood** command to control the rate at which LSA updates occur, thereby preventing high CPU or buffer utilization that can result when an area is flooded with a very large number of LSAs.

The default settings for OSPFv3 packet pacing timers are suitable for the majority of OSPFv3 deployments. Do not change the packet pacing timers unless all other options to meet OSPFv3 packet flooding requirements have been exhausted. Specifically, network operators should prefer summarization, stub area usage, queue tuning, and buffer tuning before changing the default flood timers. Furthermore, no guidelines exist for changing timer values; each OSPFv3 deployment is unique and should be considered on a case-by-case basis. The network operator assumes risks associated with changing the default flood timer values.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to configure LSA flood packet-pacing updates to occur in 55-millisecond intervals for OSPFv3 routing process 1:

```
RP/0/RSP0/CPU0:router(config)# router ospfv3 1
RP/0/RSP0/CPU0:router(config-ospfv3)# timers pacing flood 55
```

Related Commands

Command	Description
show ospfv3, on page 1370	Displays general information about OSPFv3 routing processes.
timers pacing lsa-group, on page 1429	Changes the interval at which OSPFv3 link-state advertisements (LSAs) are collected into a group and refreshed, checksummed, or aged
timers pacing retransmission , on page 1431	Configures LSA retransmission packet pacing.

timers pacing lsa-group

To change the interval at which Open Shortest Path First Version 3 (OSPFv3) link-state advertisements (LSAs) are collected into a group and refreshed, checksummed, or aged, use the **timers pacing lsa-group** command in an appropriate configuration mode. To restore the default value, use the **no** form of this command.

timers pacing lsa-group *seconds*
no timers pacing lsa-group

Syntax Description

seconds	Interval (in seconds) at which LSAs are grouped and refreshed, checksummed, or aged. Range is 10 to 1800 seconds.

Command Default

seconds : 240

OSPFv3 LSA group pacing is enabled by default.

Command Modes

Router OSPFv3 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **timers pacing lsa-group** command to control the rate at which LSA updates occur so that high CPU or buffer utilization that can occur when an area is flooded with a very large number of LSAs can be reduced. The default settings for OSPFv3 packet pacing timers are suitable for the majority of deployments. Do not change the packet pacing timers unless all other options to meet OSPFv3 packet flooding requirements have been exhausted. Specifically, network operators should prefer summarization, stub area usage, queue tuning, and buffer tuning before changing the default flooding timers. Furthermore, no guidelines exist for changing timer values; each OSPFv3 deployment is unique and should be considered on a case-by-case basis. The network operator assumes the risks associated with changing the default timer values.

Cisco IOS XR software groups the periodic refresh of LSAs to improve the LSA packing density for the refreshes in large topologies. The group timer controls the interval used for group refreshment of LSAs; however, this timer does not change the frequency that individual LSAs are refreshed (the default refresh rate is every 30 minutes).

The duration of the LSA group pacing is inversely proportional to the number of LSAs the router is handling. For example, if you have about 10,000 LSAs, decreasing the pacing interval would benefit you. If you have a very small database (40 to 100 LSAs), increasing the pacing interval to 10 to 20 minutes might benefit you slightly.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to configure OSPFv3 group packet-pacing updates between LSA groups to occur in 60-second intervals for OSPFv3 routing process 1:

```
RP/0/RSP0/CPU0:router(config)# router ospfv3 1
RP/0/RSP0/CPU0:router(config-ospfv3)# timers pacing lsa-group 60
```

Related Commands

Command	Description
show ospfv3, on page 1370	Displays general information about OSPFv3 routing processes.
timers pacing flood, on page 1427	Configures LSA flood packet pacing.
timers pacing retransmission , on page 1431	Configures LSA retransmission packet pacing.

timers pacing retransmission

To configure link-state advertisement (LSA) retransmission packet pacing, use the **timers pacing retransmission** command in an appropriate configuration mode. To restore the default retransmission packet pacing value, use the **no** form of this command.

timers pacing retransmission *milliseconds*
no timers pacing retransmission

Syntax Description

milliseconds	Time (in milliseconds) at which LSAs in the retransmission queue are paced. Range is 5 milliseconds to 100 milliseconds.

Command Default

milliseconds : 66

Command Modes

Router OSPFv3 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **timers pacing retransmission** command to control interpacket spacing between consecutive link-state update packets in the OSPFv3 retransmission queue. This command controls the rate at which LSA updates occur. When an area is flooded with a very large number of LSAs, the LSA updates can result in high CPU or buffer utilization. Using this command reduces CPU or buffer utilization.

The default settings for OSPFv3 packet retransmission pacing timers are suitable for the majority of deployments. Do not change the packet retransmission pacing timers unless all other options to meet OSPFv3 packet flooding requirements have been exhausted. Specifically, network operators should prefer summarization, stub area usage, queue tuning, and buffer tuning before changing the default flooding timers. Furthermore, no guidelines exist for changing timer values; each OSPFv3 deployment is unique and should be considered on a case-by-case basis. The network operator assumes risks associated with changing the default packet retransmission pacing timer values.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to configure LSA flood pacing updates to occur in 55-millisecond intervals for OSPFv3 routing process 1:

```
RP/0/RSP0/CPU0:router(config)# router ospfv3 1
```

```
RP/0/RSP0/CPU0:router(config-ospfv3)# timers pacing retransmission 55
```

Related Commands

Command	Description
show ospfv3, on page 1370	Displays general information about OSPFv3 routing processes.
timers pacing flood, on page 1427	Configures LSA flood packet pacing.
timers pacing lsa-group, on page 1429	Changes the interval at which OSPFv3 LSAs are collected into a group and refreshed, checksummed, or aged.

timers throttle lsa all (OSPFv3)

To set rate-limiting values for Open Shortest Path First Version 3 (OSPFv3) link-state advertisement (LSA) generation, use the **timers throttle lsa all** command in an appropriate configuration mode. To restore the default values, use the **no** form of this command.

timers throttle lsa all *start-interval hold-interval max-interval*
no timers throttle lsa all

Syntax Description

start-interval	Minimum delay (in milliseconds) for the generation of LSAs. The first instance of LSA is always generated immediately upon a local OSPFv3 topology change. The generation of the next LSA is not before the start interval. Range is 0 to 600000 milliseconds.
hold-interval	Incremental time (in milliseconds). This value is used to calculate the subsequent rate limiting times for LSA generation. Range is 1 to 600000 milliseconds.
max-interval	Maximum wait time (in milliseconds) between generation of the same LSA. Range is 1 to 600000 milliseconds.

Command Default

start-interval : 500 milliseconds

hold-interval : 5000 milliseconds

max-interval : 5000 milliseconds

Command Modes

Router OSPFv3 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

The “same LSA” is defined as an LSA instance that contains the same LSA ID number, LSA type, and advertising router ID. We recommend that you keep the *milliseconds* value of the **timers lsa arrival** command less than or equal to the *hold-interval* value of the **timers throttle lsa all** command.

Task ID

Task ID	Operations
ospf	read, write

Examples

This example shows how to customize OSPFv3 LSA throttling so that the start interval is 200 milliseconds, the hold interval is 10,000 milliseconds, and the maximum interval is 45,000 milliseconds. The minimum interval between instances of receiving the same LSA is 2000 milliseconds.

```
RP/0/RSP0/CPU0:router(config)# router ospfv3 1
RP/0/RSP0/CPU0:router(config-ospfv3)# timers throttle lsa all 200 10000 45000
RP/0/RSP0/CPU0:router(config-ospfv3)# timers lsa arrival 2000
```

Related Commands

Command	Description
show ospfv3, on page 1370	Displays general information about OSPFv3 routing processes.
timers lsa arrival, on page 1426	Sets the minimum interval at which the software accepts the same LSA from OSPFv3 neighbors.

timers throttle spf (OSPFv3)

To turn on Open Shortest Path First Version 3 (OSPFv3) shortest path first (SPF) throttling, use the **timers throttle spf** command in an appropriate configuration mode. To turn off SPF throttling, use the **no** form of this command.

timers throttle spf *spf-start spf-hold spf-max-wait*
no timers throttle spf *spf-start spf-hold spf-max-wait*

Syntax Description

spf-start	Initial SPF schedule delay (in milliseconds). Range is 1 to 600000 milliseconds.
spf-hold	Minimum hold time (in milliseconds) between two consecutive SPF calculations. Range is 1 to 600000 milliseconds.
spf-max-wait	Maximum wait time (in milliseconds) between two consecutive SPF calculations. Range is 1 to 600000 milliseconds.

Command Default

spf-start : 5000 milliseconds

spf-hold: 10000 milliseconds

spf-max-wait: 10000 milliseconds

Command Modes

Router OSPFv3 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

The first wait interval between SPF calculations is the amount of time (in milliseconds) specified by the *spf-start* argument. Each consecutive wait interval is twice the current hold level (in milliseconds) until the wait time reaches the maximum time (in milliseconds) as specified by the *spf-max-wait* argument. Subsequent wait times remain at the maximum until the values are reset or a link-state advertisement (LSA) is received between SPF calculations.



Tip Setting a low *spf-start* time and *spf-hold* time causes routing to switch to the alternate path more quickly if a failure occurs. However, it consumes more CPU processing time.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to change the start, hold, and maximum wait interval values to 5, 1000, and 90,000 milliseconds, respectively:

```
RP/0/RSP0/CPU0:router(config)# router ospfv3 1
RP/0/RSP0/CPU0:router(config-ospfv3)# timers throttle spf 5 1000 90000
```

trace (OSPFv3)

To specify the Open Shortest Path First Version 3 (OSPFv3) buffer size, use the **trace** command in router ospfv3 configuration mode. To return to the default value, use the **no** form of this command.

trace size *buffer_name size*
no trace size *buffer_name size*

Syntax Description

size	Deletes existing buffer and creates one with *N* entries.
buffer_name	Specifies a buffer from one of the 15 listed buffers. Refer Table 142: Buffer Types, on page 1437 table for details on the buffers.
size	Specifies allowed size for the selected buffer. Options are: 0, 256, 1024, 2048, 4096, 8192, 16384, 32768, and 65536. Select 0 to disable traces.

Command Default

No default behavior or values

Command Modes

Router ospfv3 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Trace buffers are used to store various traffic and processing events during the runtime. Large buffers can store more events. If the buffer becomes full, old entries are overwritten by the latest entries. In a large network, user may want to increase the trace buffer size to accommodate more events.

Table 142: Buffer Types

Name	Description
adj	adjacency
adj_cycle	dbd/flood events/pkts
config	config events
errors	errors
events	mda/rtrid/bfd/vrf
ha	startup/HA/NSF
hello	hello events/pkts
idb	interface

Name	Description
pkt	I/O packets
rib	rib batching
spf	spf/topology
spf_cycle	spf/topology detail
te	mpls-te
test	testing info
mq	message queue info

Task ID

Task ID	Operations
ospf	read, write

Examples

This example shows how to set 1024 error trace entries:

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#router ospfv3 osp3
RP/0/RSP0/CPU0:router(config-ospfv3)#trace size errors ?
  0       disable trace
  256     trace entries
  512     trace entries
  1024    trace entries
  2048    trace entries
  4096    trace entries
  8192    trace entries
  16384   trace entries
  32768   trace entries
  65536   trace entries
RP/0/RSP0/CPU0:router(config-ospfv3)#trace size errors 1024
```

transmit-delay (OSPFv3)

To set the estimated time required to send a link-state update packet on the interface, use the **transmit-delay** command in an appropriate configuration mode. To return to the default value, use the **no** form of this command.

transmit-delay *seconds*
no transmit-delay *seconds*

Syntax Description

seconds	Time (in seconds) required to send a link-state update. Range is 1 to 65535 seconds.

Command Default

1 second

Command Modes

Process configuration

Area configuration

Interface configuration

Virtual-link 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Link-state advertisements (LSAs) in the update packet must have their ages incremented by the amount specified in the *seconds* argument before transmission. The value assigned should take into account the transmission and propagation delays for the interface.

If the delay is not added before transmission over a link, the time in which the LSA propagates over the link is not considered. This setting has more significance on very low-speed links.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to configure a transmit delay of 3 seconds for GigabitEthernet interface 0/3/0/0:

```
RP/0/RSP0/CPU0:router(config)# router ospfv3 1
RP/0/RSP0/CPU0:router(config-ospfv3)# area 0
RP/0/RSP0/CPU0:router(config-ospfv3-ar)# interface GigabitEthernet 0/3/0/0
RP/0/RSP0/CPU0:router(config-ospfv3-ar-if)# transmit-delay 3
```

Related Commands

Command	Description
show ospfv3, on page 1370	Displays general information about OSPF routing processes.

virtual-link (OSPFv3)

To define an Open Shortest Path First Version 3 (OSPFv3) virtual link, use the **virtual-link** command in area configuration mode. To remove a virtual link, use the **no** form of this command.

virtual-link *router-id*
no virtual-link

Syntax Description

router-id	Router ID associated with the virtual link neighbor. The router ID appears in the **show ospfv3** display. This value must be entered in 32-bit dotted-decimal notation, similar to an IP Version 4 (IPv4) address. There is no default.

Command Default

No virtual links are defined.

Command Modes

Area 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

In OSPFv3, when there exists a path through another non-backbone area over which the virtual link can function, all areas must be connected to a backbone area. If the connection to the backbone is lost, it can be repaired by establishing a virtual link.

Virtual links, which are defined in the submode of the area they transit, are in effect virtual point-to-point interfaces belonging to area 0 (the backbone). The virtual links inherit parameter values from the backbone area, rather than the transit area in which they are defined.

Each virtual link neighbor must include the router ID of the virtual link neighbor for the link to be properly established. Use the **show ospfv3** command to display the router ID of an OSPFv3 process.

Use the **virtual-link** command to place the router in virtual-link configuration mode (config-router-ar-vl), from which you can configure virtual-link-specific settings. Commands configured under this mode (such as the **transmit-delay** command) are automatically bound to that virtual link.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to establish a virtual link with default values for all optional parameters:

```
RP/0/RSP0/CPU0:router(config)# router ospfv3 201
```

```
RP/0/RSP0/CPU0:router(config-ospfv3)# area 1
RP/0/RSP0/CPU0:router(config-ospfv3-ar)# virtual-link 10.3.4.5
```

Related Commands

Command	Description
show ospfv3, on page 1370	Displays general information about OSPF routing processes.
transmit-delay (OSPFv3), on page 1439	Sets the estimated time required to send a link-state update packet on the interface.

vrf (OSPFv3)

To configure an Open Shortest Path First Version 3 (OSPFv3) VPN routing and forwarding (VRF) instance and enter OSPFv3 VRF configuration submode, use the **vrf** command in router configuration mode. To terminate an OSPFv3 VRF, use the **no** form of this command.

vrf *vrf-name*
no vrf *vrf-name*

Syntax Description

vrf-name	Specifies a name for the OSPFV3 vrf. If a name is not specified, the default vrf is assumed.

Command Default

No OSPFv3 VRF is configured.

Command Modes

Router configuration

Command History

Release	Modification
Release 4.1.0	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.

Use the **vrf** command to explicitly configure a VRF. This command creates a separate instance of OSPFv3 within the process. Commands configured under the VRF configuration mode (such as the **interface** [OSPFv3] and **authentication** [OSPFv3] commands) are automatically bound to that VRF.

To modify or remove the VRF, the *vrf-name* argument format must be the same as the format used when creating the VRF.

Task ID

Task ID	Operation
ospf	read, write

This example shows how to configure VRF *vrf_1* and enter OSPFv3 VRF configuration submode:

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#router ospfv3 osp3
RP/0/RSP0/CPU0:router(config-ospfv3)#vrf vrf_1
RP/0/RSP0/CPU0:router(config-ospfv3-vrf)#
```

Related Commands

Command	Description
router-id (OSPFv3), on page 1365	Configures a router ID for an OSPFv3 process.



RIB Commands

This module describes the commands used to display and clear information in the Routing Information Base (RIB).

For detailed information about RIB concepts, configuration tasks, and examples, see the *Implementing RIB on* Cisco ASR 9000 Series Router module in *Routing Configuration Guide for Cisco ASR 9000 Series Routers*.

- address-family next-hop dampening disable, on page 1447
- clear route, on page 1448
- maximum prefix (RIB), on page 1450
- lcc, on page 1452
- rcc, on page 1453
- recursion-depth-max, on page 1455
- router rib, on page 1456
- rump always-replicate, on page 1457
- show lcc statistics, on page 1458
- show mpls lsd forwarding, on page 1460
- show rcc, on page 1461
- show rcc statistics, on page 1463
- show rcc vrf, on page 1465
- show rib, on page 1466
- show rib afi-all, on page 1468
- show rib attributes, on page 1470
- show rib client-id, on page 1471
- show rib clients, on page 1473
- show rib extcomms, on page 1475
- show rib firsthop, on page 1477
- show rib history, on page 1479
- show rib next-hop, on page 1481
- show rib opaques, on page 1483
- show rib protocols, on page 1485
- show rib recursion-depth-max, on page 1487
- show rib statistics, on page 1489
- show rib tables, on page 1492
- show rib trace, on page 1494

- show rib vpn-attributes, on page 1496
- show rib vrf, on page 1498
- show route, on page 1500
- show route backup, on page 1506
- show route best-local, on page 1509
- show route connected, on page 1511
- show route local, on page 1513
- show route longer-prefixes, on page 1515
- show route next-hop, on page 1517
- show route quarantined, on page 1519
- show route resolving-next-hop, on page 1521
- show route static, on page 1523
- show route summary, on page 1525

address-family next-hop dampening disable

To disable Routing Information Base (RIB) next-hop dampening, use the **address-family next-hop dampening disable** command in router configuration mode. To enable RIB next-hop dampening, use the **no** form of this command.

address-family {**ipv4** | **ipv6**} **next-hop dampening disable**
no address-family {**ipv4** | **ipv6**} **next-hop dampening disable**

Syntax Description

ipv4	Specifies IP Version 4 (IPv4) address prefixes.
ipv6	Specifies IP Version 6 (IPv6) address prefixes.

Command Default

RIB next-hop dampening is enabled.

Command Modes

Router 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 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
rib	read, write

Examples

The following example shows how to disable RIB next-hop dampening for IPv6 address families:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router rib
RP/0/RSP0/CPU0:router(config-rib)# address-family ipv6 next-hop dampening disable
```

clear route

To clear routes from the IP routing table, use the **clear route** command in EXEC configuration mode.

clear route [**vrf** {*vrf-name* | **all**}] {**ipv4** | **ipv6** | **afi-all**} {**unicast** | **multicast** | **safi-all**} [**topology** *topo-name*] [*ip-address mask*]

Syntax Description

vrf { *vrf-name* | **all** }	(Optional) Specifies a particular VPN routing and forwarding (VRF) instance or all VRF instances.
ipv4	Specifies IP Version 4 address prefixes.
ipv6	Specifies IP Version 6 address prefixes.
afi -all	Specifies IP Version 4 and IP Version 6 address prefixes.
unicast	Specifies unicast address prefixes.
multicast	Specifies multicast address prefixes.
safi-all	Specifies unicast and multicast address prefixes.
topology *topo-name*	(Optional) Specifies topology table information and name of the topology table.
ip-address node-id	(Optional) Clears hardware resource counters from the designated node. The *node-id* argument is entered in the *rack/slot/module* notation.
ip-address	Network IP address about which routing information should be displayed.
mask	Network mask specified in either of two ways: Network mask can be a four-part, dotted-decimal address. For example, 255.0.0.0 indicates that each bit equal to 1 means the corresponding address bit is a network address. Network mask can be indicated as a slash (/) and number. For example, /8 indicates that the first 8 bits of the mask are 1s, and the corresponding bits of the address are the network address.

Command Default

If a **vrf** *vrf-name* is not specified, routes are cleared from the default IPv4 unicast VRF.

Command Modes

EXEC 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **clear route** command to clear routes from an IP routing table to a specific network, a matching subnet address, or all routes.

Task ID

Task ID	Operations
rib	read, write

Examples

The following example shows how to remove all routes matching the subnet address 192.168.2.0 and mask 255.255.255.0 from the IPv4 unicast routing table:

```
RP/0/RSP0/CPU0:router# clear route ipv4 unicast 192.168.2.0 255.255.255.0
```

The following example shows how to remove all routes from the IPv4 unicast routing table:

```
RP/0/RSP0/CPU0:router# clear route ipv4 unicast
```

Related Commands

Command	Description
show route, on page 1500	Displays the current state of the routing table.

maximum prefix (RIB)

To set the prefix limit for the VPN routing and forwarding (VRF) instance, use the **maximum prefix** command in global VRF address family configuration mode. To set the prefix limits to the default values, use the **no** form of this command.

maximum prefix *maximum* [*mid-threshold*]
no maximum prefix

Syntax Description

maximum	Maximum number of prefixes allowed in the VRF instance. Range is 32 to 2000000.
mid-threshold	(Optional) Integer specifying at what percentage of the *maximum* argument value the software starts to generate a Simple Network Management Protocol (SNMP) trap. Range is 1 to 100.

Command Default

No default behavior or values

Command Modes

Global VRF address family configuration

Command History

Release	Modification
Release 3.7.2	This command was introduced.
Release 6.5.3	This command was modified. Routes will be blocked after the set mid-threshold value is reached.

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.

You can set a value for the *mid-threshold* argument which is a hard limit. Routes will be blocked after the set mid-threshold value is reached.

Use the **maximum prefix** command to configure a maximum number of prefixes that a VRF instance is allowed to receive.

Task ID

Task ID	Operations
rib	read, write

Examples

The following example shows how to set the maximum number of prefixes allowed to 1000:

```
RP/0/RSP0/CPU0:router(config)# vrf vrf-A
RP/0/RSP0/CPU0:router(config-vrf)# address-family ipv4 unicast
RP/0/RSP0/CPU0:router(config-vrf-af)#  maximum prefix 1000
```

A maximum number of routes is applicable to dynamic routing protocols as well as static or connected routes. When maximum prefix is configured, an **syslog** message is generated in the following conditions:

1. if the number of routes is above the threshold when“maximum prefix” configuration has been committed
2. if the number routes reaches the configured “maximum prefix” values for the VRF

Related Commands

Command	Description
show rib tables, on page 1492	Displays all tables known to the RIB.

lcc

To enable Label Consistency Checker (lcc) background scan for IPv6 or IPv4 labels, use the **lcc enable** command in global configuration mode. To disable lcc background scan, use the **no** for of this command.

lcc { **ipv4** | **ipv6** } **unicast** { **enable** | **period** *milliseconds* }
no lcc { **ipv4** | **ipv6** } **unicast** { **enable** | **period** *milliseconds* }

Syntax Description

ipv4	Specifies IP Version 4 address prefixes.
ipv6	Specifies IP Version 6 address prefixes.
unicast	Specifies unicast address prefixes.
period *milliseconds*	Specifies the period between scans in milliseconds. Range is 100 to 600000 milliseconds.

Command Default

Label consistency checker is disabled.

Command Modes

Global configuration

Command History

Release	Modification
Release 4.2.0	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
ipv4	read, write
ipv6	read, write

This example shows how to enbale lcc for IPv6 lables:

```
Router#configure
Router(config)#lcc ipv6 unicast enable
```

rcc

To ensure consistency between centralized and distributed routing information bases (RIBs), use the **rcc** command in global configuration mode. To disable this function, use the **no** form of this command.

rcc { **ipv4** | **ipv6** } **unicast** { **enable** | **period** *milliseconds* }
no rcc { **ipv4** | **ipv6** } **unicast** { **enable** | **period** *milliseconds* }

Syntax Description

ipv4	Specifies IP Version 4 address prefixes.
ipv6	Specifies IP Version 6 address prefixes.
unicast	Specifies unicast address prefixes.
period *milliseconds*	Specifies the period between scans in milliseconds. Range is 100 to 600000 milliseconds.

Command Default

Route consistency checker is disabled.

Command Modes

Global configuration

Command History

Release	Modification
Release 4.2.0	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.

Use the **period** option to control how often the scan be triggered. Each time the scan is triggered, the background scan process resumes verification from where it was left out and sends one buffer's worth of routes to the forwarding information base (FIB).

Task ID

Task ID	Operation
ipv4	read, write
ipv6	read, write

This example shows how to configure rcc for IPv6 unicast:

```
Router#configure
Router(config)#rcc ipv6 unicast enable
```

This example shows how to enable rcc with a scan period of 500 milliseconds for IPv6 unicast:

```
Router#configure
Router(config)#rcc ipv6 unicast period 500
```

recursion-depth-max

To set the maximum depth for route recursion checks, use the **recursion-depth-max** command in router configuration mode. To set the recursion checks to the default value, use the **no** form of this command.

recursion-depth-max *maximum*
no recursion-depth-max *maximum*

Syntax Description

maximum	Maximum depth for recursion checks. Range is 5 to 16.

Command Default

The default recursion depth is 128.

Command Modes

Router 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **recursion-depth-max** command to configure a specific maximum number of recursion checks in the range of 5 to 16.

Task ID

Task ID	Operations
rib	read, write

Examples

The following example shows how to set the maximum depth for route recursion checks to 12:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router rib
RP/0/RSP0/CPU0:router(config-rib)# recursion-depth-max 12
```

router rib

To enter Routing Information Base (RIB) configuration mode, use the **router rib** command in global configuration mode. To remove all RIB configurations and terminate the RIB routing process, use the **no** form of this command.

router rib
no router rib

Syntax Description

This command has no arguments or keywords.

Command Default

Router configuration mode is not enabled.

Command Modes

Global 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 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
ospf	read, write
hsrp	read, write
isis	read, write

Examples

The following example shows how to enter RIB configuration mode:

```
RP/0/RSP0/CPU0:router(config)# router rib
```

rump always-replicate

To enable replication from uRIB to muRIB as usual even after features such as MTR are configured, use the **rump always-replicate** command in router configuration mode. To diable replication from uRIB to muRIB, use the **no** form of this command.

rump always-replicate [*access-list*]
no rump always-replicate [*access-list*]

Syntax Description

access-list-name	(Optional) Name of the access list.

Command Default

Replication from uRIB to muRIB is enabled.

Command Modes

Router address family configuration

Command History

Release	Modification
Release 3.9.0	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.

Configuring the **rump always-replicate** command allows routers in a network to be upgraded to multitopology routing gradually without a flag day where all routers need to be configured at the same time without major service disruption. When **rump always-replicate** is configured, replicated routes are added into the muRIB with the lowest admin distance. So if protocols are populating the muRIB, they continue to do so. For the same route, protocol routes win over replicated routes because of higher admin distance.

If an unwanted more specific route comes from the uRIB, optionally provide an access list through which the replicated routes are run. If the route passes the access list, the route is replicated by RUMP.

Task ID

Task ID	Operations
rib	read, write

Examples

The following example shows how to enale replication from uRIB to muRIB:

```
RP/0/RSP0/CPU0:router(config)# router rib
RP/0/RSP0/CPU0:router(config-rib)# address-family ipv4
RP/0/RSP0/CPU0:router(config-rib-afi)# rump always-replicate
```

show lcc statistics

To view results of a label consistency checker (lcc) background scan, use the **show lcc statistics** command in EXEC mode.

show lcc {**ipv4** | **ipv6**} **unicast statistics** {**scan-id** | **summary**}

Syntax Description

ipv4		IPv4 address prefix.
ipv6		IPv6 address prefix.
unicast		Specifies unicast address prefix.
scan-id	*scan-id-value*	Specifies the scan ID value. The range is between <0-100000>.
summary		Displays a summary of the BG route consistency check statistics.

Command Default

None

Command Modes

EXEC

Command History

Release	Modification
Release 4.2.0	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
ipv4	read
ipv6	read

This example shows background scan statistics for AFI-SAFI mplsv6-unicast:

```
RP/0/RSP0/CPU0:router#show lcc ipv6 unicast statistics

Background Scan Statistics for AFI-SAFI mplsv6-unicast:
=======================================================

Scan enabled:           False
Current scan-id:        0               Scan triggered:        False
Configured period:      60              Current period:        0

Paused by range scan: False
Paused by route churn: False
Paused by error scan: False
```

```
Last data sent: 0 entries                Damping percent:       70
Default route churn:    10               Current route churn:   0
Route churn last calculated at           Dec 31 16:00:00.000

Logs stored for background scan ids:

Log for AFI-SAFI mplsv6-unicast:
================================

End Of Logs
```

This example shows background scan statistics for AFI-SAFI mplsv4-unicast:

```
RP/0/RSP0/CPU0:router#show lcc ipv4 unicast statistics


Background Scan Statistics for AFI-SAFI mplsv4-unicast:
=======================================================

Scan enabled:           False
Current scan-id:        0                Scan triggered:        False
Configured period:      60               Current period:        0

Paused by range scan: False
Paused by route churn: False
Paused by error scan: False

Last data sent: 0 entries                Damping percent:       70
Default route churn:    10               Current route churn:   0
Route churn last calculated at           Dec 31 16:00:00.000

Logs stored for background scan ids:

Log for AFI-SAFI mplsv4-unicast:
================================

End Of Logs
```

show mpls lsd forwarding

To show the priority and version information of an LSD rewrite, use the **show mpls lsd forwarding** command with **detail** option.

show mpls lsd forwarding detail

Syntax Description

detail	Use the **detail** keyword to show the priority and version information of an LSD rewrite.

Command Default

None

Command Modes

EXEC

Command History

Release	Modification
Release 4.2.0	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
mpls	read

This example shows the output from **show mpls lsd forwarding detail** command:

```
RP/0/RSP0/CPU0:router# show mpls lsd forwarding

In_Label, (ID), Path_Info: <Type>
16002, (IPv4, 'default':4U, 2.1.1.1/32), 1 Paths
   1/1: IPv4, 'default':4U, Gi0/3/0/3, nh=20.1.5.2, lbl=None, tun_id=0 flags=()
            path-id=1, backup-path-id=0
      BCDL priority:3, LSD queue:5, version:7
      Installed Jan 17 14:26:31.140 (05:03:11 ago)
```

This table describes the significant fields shown in the display

Field	Description
BCDL priority:	BCDL priority value
LSD queue:	LSD queue number
version:	Version number of the LSD rewrite

show rcc

To display route consistency checker (RCC) information, use the **show rcc** command in EXEC mode.

show rcc {**ipv4** | **ipv6**} **unicast** [*prefix netmask* **vrf** *vrf-name*]

Syntax Description

ipv4	Specifies IP Version 4 address prefixes.
ipv6	Specifies IP Version 6 address prefixes.
unicast	Specifies unicast address prefixes.
prefix	(Optional) Starting prefix.
netmask	(Optional) Network mask.
vrf *vrf-name*	(Optional) Specifies a particular VPN routing and forwarding (VRF) instance or all VRF instances.

Command Default

No default behavior or values

Command Modes

EXEC

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 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
ipv4	read

Examples

The following is sample output from the **show rcc** command:

```
  RP/0/RSP0/CPU0:router# show rcc ipv4 unicast statistics
Thu Mar 26 13:47:28.391 IST

Background Scan Summary
=======================

Scan enabled:          False            Last scan-id:  0
Configured period:     15000            Current period:        0

Paused By:
  route churn:False  on-demand scan:False  error scan:False
```

```
Last data sent: 0 entries                Damping percent:     69
Default route churn:    100              Current route churn: 0
Route churn last calculated at           Never
Logs last cleared at                     Never

Scan paused by ISSU                      False

Logs stored for background scan ids:

Scan Logs
=========
Legend:
        ? - Currently Inactive Node, ! - Non-standard SVD Role
        * - Node did not reply

End of Logs
```

show rcc statistics

To view results of a route consistency checker (rcc) background scan, use the **show rcc statistics** command in EXEC mode.

show rcc {**ipv4** | **ipv6**} **unicast statistics** {**scan-id** | **summary**}

Syntax Description

ipv4	IPv4 address prefix.
ipv6	IPv6 address prefix.
unicast	Specifies unicast address prefixes.
scan-id *scan-id-value*	Specifies the scan ID value. The range is between <0-100000>.
summary	Displays a summary of the BG route consistency check statistics.

Command Default

None

Command Modes

EXEC

Command History

Release	Modification
Release 4.2.0	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
ipv4	read
ipv6	read

This example shows background scan statistics for AFI-SAFI IPv6 unicast:

```
RP/0/RSP0/CPU0:router#show rcc ipv6 unicast statistics

Background Scan Statistics for AFI-SAFI ipv6-unicast:
=====================================================

Scan enabled:           False
Current scan-id:        0               Scan triggered:        False
Configured period:      60              Current period:        0

Paused by range scan: False
Paused by route churn: False
Paused by error scan: False
```

```
Last data sent: 0 entries                Damping percent:       70
Default route churn:    10               Current route churn:   0
Route churn last calculated at           Dec 31 16:00:00.000

Logs stored for background scan ids:

Log for AFI-SAFI ipv6-unicast:
================================

End Of Logs
```

This example shows background scan statistics for AFI-SAFI Ipv4 unicast:

```
RP/0/RSP0/CPU0:router#show rcc ipv4 unicast statistics

Background Scan Statistics for AFI-SAFI ipv4-unicast:
=====================================================

Scan enabled:           False
Current scan-id:        0                Scan triggered:        False
Configured period:      60               Current period:        0

Paused by range scan: False
Paused by route churn: False
Paused by error scan: False

Last data sent: 0 entries                Damping percent:       70
Default route churn:    10               Current route churn:   0
Route churn last calculated at           Dec 31 16:00:00.000

Logs stored for background scan ids:

Log for AFI-SAFI ipv4-unicast:
================================

End Of Logs
```

show rcc vrf

To run on-demand route consistency checker (rcc) scan on AFI, SAFI, table, and prefix or the entire set of prefixes in the table, use the **show rcc vrf** command in EXEC mode.

show rcc {**ipv4** | **ipv6**} **unicast** *prefix/mask* **vrf** *vrfname*

Syntax Description

ipv4	IPv4 address prefix.
ipv6	IPv6 address prefix.
prefix /mask	Specifies unicast address prefix.
vrf	Specifies VPN routing and forwarding (VRF) instance.
vrfname	Name of the VRF.

Command Default

None.

Command Modes

EXEC

Command History

Release	Modification
Release 4.2.0	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
ipv4	read
ipv6	read

This example shows how to run on-demand rcc scan for an IPv6 prefix:

```
RP/0/RSP0/CPU0:router#show rcc ipv6 unicast 2001:DB8::/32 vrf vrf_1
```

This example shows how to run on-demand rcc scan for an Ipv4 prefix:

```
RP/0/RSP0/CPU0:router#show rcc ipv4 unicast 10.2.3.4/32 vrf vrf-1
```

show rib

To display Routing Information Base (RIB) data, use the **show rib** command in EXEC mode.

show rib {**ipv4** | **ipv6**} {**unicast** | **multicast**}[**firsthop** | [*type interface-path-id*] | **next-hop** | [*type interface-path-id*] | **opaques** | {**attribute** | **ip-nexthop** | **ipfrr** | **safi-tunnel** | **summary** | **tunnel-nexthop**} | **protocols** | [**standby**] | **statistics** | [*name*] | [**standby**] | **topology** | {*topo-name* | **all**}]

Syntax Description

ipv4	(Optional) Specifies IP Version 4 address prefixes.
ipv6	(Optional) Specifies IP Version 6 address prefixes.
unicast	(Optional) Specifies unicast address prefixes. This is the default.
multicast	(Optional) Specifies multicast address prefixes.
firsthop	(Optional) Specifies registered first-hop notification addresses.
type	Interface type. For more information, use the question mark (?) online help function.
interface-path-id	Identifies a physical interface or a virtual interface. **Note** Use the show interfaces command to see a list of all possible interfaces currently configured on the router. For more information about the syntax for the router, use the question mark (?) online help function.
next-hop	(Optional) Specifies registered next-hop notification addresses.
opaques	(Optional) Specifies opaque data installed in the RIB.
attribute	(Optional) Specifies opaque attributes installed in the RIB.
ip-nexthop	(Optional) Specifies P next-hop data installed in the RIB.
safi-tunnel	(Optional) Specifies subaddress family (SAFI) tunnel opaque data installed in the RIB.
summary	(Optional) Specifies a summary of opaque data installed in the RIB.
tunnel-nexthop	(Optional) Specifies tunnel next-hop opaque data installed in the RIB.
protocols	(Optional) Specifies registered protocols.
statistics *name*	(Optional) Specifies RIB statistics of a given name.
standby	(Optional) Specifies standby information.
all	(Optional) Specifies that all topology table information should be displayed.

Command Default

No default behavior or values

Command Modes

EXEC

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 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
ipv4	read

Examples

The following example illustrates the **show rib** command:

```
RP/0/RSP0/CPU0:router# show rib

            ipv4 multicast


            topology
          BLUE

RP/0/RSP0/CPU0:router# show rib topology BLUE ipv4 multicast protocols
Protocol   Handle   Instance
isis         0      mt
```

Related Commands

Command	Description
show rib afi-all, on page 1468	Displays both IPv4 and IPv6 RIB information.

show rib afi-all

To display Routing Information Base (RIB) data for both IPv4 and IPv6 address families, use the **show rib afi-all** command in EXEC mode.

show rib afi-all [**attributes**] [**client-id**] [**clients**] [**extcomms**] [**firsthop**] [**history**] [**multicast**] [**next-hop**] [**opaques**] [**protocols**] [**recursion-depth-max**] [**safi-all**] [**statistics**] [**tables**] [**trace**] [**unicast**] [**vpn-attributes**]

Syntax Description

attributes	(Optional) Displays all BGP attributes installed in RIB.
client-id	(Optional) Displays RIB client ID for longer history of redistributed routes sent to the client.
clients	(Optional) Displays RIB clients.
extcomms	(Optional) Displays all extended communities installed in RIB.
firsthop	(Optional) Displays registered firsthop notification addresses.
history	(Optional) Displays redistributed routes sent to RIB clients.
multicast	(Optional) Displays multicast commands.
next-hop	(Optional) Displays registered next-hop notification addresses.
opaques	(Optional) Displays opaquae data installed in RIB.
protocols	(Optional) Displays registered protocols.
recursion-depth-max	(Optional) Displays maximum recursion depth in RIB.
safi-all	(Optional) Displays unicast and multicast commands.
statistics	(Optional) Displays RIB statistics.
tables	(Optional) Displays a list of tables known to RIB.
trace	(Optional) Displays RIB trace entries.
unicast	(Optional) Displays unicast commands.
vpn-attributes	(Optional) Displays all VPN attributes installed in RIB.

Command Default

No default behavior or values

Command Modes

EXEC

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 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
ipv4	read

Examples

The following example illustrates the **show rib afi-all attributes** command:

```
RP/0/RSP0/CPU0:router# show rib afi-all attributes

BGP attribute data in IPv4 RIB:

0 Attributes, for a total of 0 bytes.

BGP attribute data in IPv6 RIB:

0 Attributes, for a total of 0 bytes.
```

Related Commands

Command	Description
show rib, on page 1466	Displays RIB information.

show rib attributes

To display Border Gateway Protocol (BGP) attributes installed in the Routing Information Base (RIB), use the **show rib attributes** command in EXEC mode.

show rib attributes [**summary**] [**standby**]

Syntax Description

summary	(Optional) Displays a summary of BGP attribute data installed in the RIB.
standby	(Optional) Displays standby information.

Command Default

No default behavior or values

Command Modes

EXEC

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 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
rib	read

Examples

The following is sample output from the **show rib attributes** command:

```
RP/0/RSP0/CPU0:router# show rib attributes

BGP attribute data in IPv4 RIB:

Attribute ID (0x2):size (68)
Attribute ID (0x3):size (52)
Attribute ID (0x4):size (68)
Attribute ID (0x5):size (52)

4 Attributes, for a total of 240 bytes.

Attribute ID : ID assigned for the attribute by BGP
size : size of the attribute data.
```

show rib client-id

To display Routing Information Base (RIB) redistribution histories, use the **show rib client-id** command in EXEC mode.

show rib client-id *id* **redistribution history** [**standby**]

Syntax Description

id	ID of the client. Range is 0 to 4294967295.
redistribution history	Displays longer history of redistributed routes sent to RIB clients.
standby	(Optional) Displays standby information.

Command Default

No default behavior or values

Command Modes

EXEC

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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **show rib client-id** command to display a history of the route additions, deletions, and updates sent from RIB to the client across VRFs.

Task ID

Task ID	Operations
rib	read

Examples

The following is sample output from the **show rib client-id** command:

```
RP/0/RSP0/CPU0:router# show rib client-id 13 redistribution history

PID     JID   Client            Location
151630  113   bcdl_agent        node0_5_CPU0
  Table ID: 0xe0000000
     S 80.80.80.0/24[1/0]           update, 5 path(s), 0x0   Jan 31 09:54:57.224
     S 80.80.80.0/24[1/0]           update, 6 path(s), 0x0   Jan 31 09:53:39.736
     S 140.140.140.0/24[1/0]        update, 1 path(s), 0x0   Jan 31 09:53:39.729
     S 80.80.80.0/24[1/0]           update, 5 path(s), 0x0   Jan 30 22:08:38.551
     S 140.140.140.0/24             deleted,                 Jan 30 22:08:38.543
     S 80.80.80.0/24[1/0]           update, 6 path(s), 0x0   Jan 30 22:03:05.889
     S 100.100.100.0/24[1/0]        update, 1 path(s), 0x0   Jan 30 22:03:05.880
```

This table describes the significant fields shown in the display.

Table 143: show rib client-id Field Descriptions

Field	Description
PID	Process ID of the client.
JID	Job ID of the client.
Client	Client name.
Location	Location node on which the client is present.

Related Commands

Command	Description
show rib clients, on page 1473	Displays RIB clients.

show rib clients

To display Routing Information Base (RIB) clients, use the **show rib clients** command in EXEC mode.

show rib [**afi-all** | **ipv4** | **ipv6**] **clients** [**protocols** | **redistribution** [**history**]] [**standby**]

Syntax Description

afi-all	(Optional) Specifies all address families.
ipv4	(Optional) Specifies IP Version 4 address prefixes. This is the default.
ipv6	(Optional) Specifies IP Version 6 address prefixes.
protocols	(Optional) Specifies client protocols.
redistribution	(Optional) Specifies protocols redistributed by clients
history	(Optional) Specifies redistributed routes sent to RIB clients.
standby	(Optional) Displays standby information.

Command Default

No default behavior or values

Command Modes

EXEC

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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **show rib clients** command to display the list of clients who have registered with RIB, what protocol routes they are redistributing, and a history of the routes sent to the client.

The maximum number of redistribution entries is 5000 for Bulk Content Downloader (BCDL) and 500 for other protocols.

Task ID

Task ID	Operations
rib	read

Examples

The following is sample output from the **show rib clients** command:

```
RP/0/RSP0/CPU0:router# show rib clients

Process             Location         Client ID  Redist   Proto
isis                node0_5_CPU0     0           insync   insync
ospf                node0_5_CPU0       2         insync   insync
```

```
RP/0/RSP0/CPU0:router# show rib clients redistribution

isis node0_5_CPU0
   ipv4  uni        vrf default   insync         route
      static                      insync
ospf node0_5_CPU0
   ipv4  uni        vrf default   insync         route
      static                      insync
      local                       insync
bgp node0_5_CPU0
   ipv4  uni        vrf abc       insync         route
      static                      insync
bcdl_agent node0_5_CPU0
   ipv4  uni        vrf default   insync         rib_fib
   ipv4  uni        vrf bar       insync         rib_fib
   ipv4  uni        vrf abc       insync         rib_fib
   ipv4  uni        vrf test      insync         rib_fib
```

This table describes the significant fields shown in the display.

Table 144: show rib clients Field Descriptions

Field	Description
Process	Client process name.
Location	Location where the client process in running.
Client ID	ID assigned to the client by RIB.
Redist	Whether the client is redistributing any protocols or not and whether it has read all routes from RIB or not. • insync—read • outsync—not read.
Proto	Whether the protocol has sent all its routes to RIB and signaled update complete or not. • insync—read • outsync—not read.

show rib extcomms

To display all extended communities installed in the Routing Information Base (RIB), use the **show rib extcomms** command in EXEC mode.

show rib [**afi-all** | **ipv4** | **ipv6**] **extcomms** [**summary**] [**standby**]

Syntax Description

afi-all	(Optional) Specifies all address families.
ipv4	(Optional) Specifies IP Version 4 address prefixes. This is the default.
ipv6	(Optional) Specifies IP Version 6 address prefixes.
summary	(Optional) Specifies a summary of all extended communities in the RIB.
standby	(Optional) Displays standby information.

Command Default

No default behavior or values

Command Modes

EXEC

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 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
rib	read

Examples

The following is sample output from the **show rib extcomms** command:

```
RP/0/RSP0/CPU0:router# show rib extcomms

Extended community data in RIB:

Extended community                          Ref count
COST:128:128:41984                             1
EIGRP route-info:0x8000:0                      1
EIGRP AD:1:25600                               1
EIGRP RHB:255:0:16384                          1
EIGRP LM:0x0:1:4470                            1
```

This table describes the significant fields shown in the display.

Table 145: show rib extcomms Field Descriptions

Field	Description
Extended Community	Type of extended communities. Different protocols can add different extended communities.
Ref Count	Number of routes referring to the Extended community.

show rib firsthop

To display registered first-hop notification addresses, use the **show rib firsthop** command in EXEC mode.

show rib [**vrf** {*vrf-name* | **all**}] [**afi-all** | **ipv4** | **ipv6**] [**unicast** | **multicast** | **safi-all**] **firsthop** [*client-name*] [*type interface-path-id* | *ip-address /prefix-length* | *ip-address mask* | **resolved** | **unresolved** | **damped**] [**summary**] [**standby**]

Syntax Description

vrf { *vrf-name* | **all** }	(Optional) Specifies a particular VPN routing and forwarding (VRF) instance or all VRF instances.
afi-all	(Optional) Specifies all address families.
ipv4	(Optional) Specifies IP Version 4 address prefixes. This is the default.
ipv6	(Optional) Specifies IP Version 6 address prefixes.
unicast	(Optional) Specifies unicast address prefixes. This is the default.
multicast	(Optional) Specifies multicast address prefixes.
safi-all	(Optional) Specifies unicast and multicast address prefixes.
client-name	(Optional) Name of the RIB client.
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.
ip-address	(Optional) Network that BGP advertises.
/ *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.
ip-address mask	(Optional) Network mask applied to the *ip-address* argument.
resolved	(Optional) Specifies resolved next-hops.
unresolved	(Optional) Specifies unresolved next-hops.
damped	(Optional) Specifies next-hops that are damped.
summary	(Optional) Specifies a summary of the next-hop information.
standby	(Optional) Displays standby information.

Command Default

If a **vrf** *vrf-name* is not specified, the registered first-hop notifications addresses are displayed for the default IPv4 unicast VRF.

Command Modes

EXEC

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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **show rib firsthop** command to display the list of first hops registered by various clients with RIB and the address and interface through which they are resolved.

Task ID

Task ID	Operations
rib	read

Examples

The following is sample output from the **show rib firsthop** command:

```
RP/0/RSP0/CPU0:router# show rib firsthop

Registered firsthop notifications:
0.0.0.0/0 via 1.1.0.1 - MgmtEth0/5/CPU0/0, ospf/node0_5_CPU0
1.1.0.1/32 via 1.1.0.1 - MgmtEth0/5/CPU0/0, ipv4_static/node0_5_CPU0
1.1.1.1/32 via 1.1.1.1 - MgmtEth0/5/CPU0/0, ipv4_static/node0_5_CPU0
10.10.10.1/32 via 10.10.10.1 - Loopback0, ipv4_static/node0_5_CPU0
10.10.10.3/32 via 10.10.10.3 - Loopback0, ipv4_static/node0_5_CPU0
15.15.15.1/32 via 10.10.10.1 - Loopback0, ipv4_static/node0_5_CPU0
20.20.20.1/32 via 1.1.1.1 - MgmtEth0/5/CPU0/0, ipv4_static/node0_5_CPU0
30.30.30.1/32 via 1.1.1.2 - MgmtEth0/5/CPU0/0, ipv4_static/node0_5_CPU0
```

show rib history

To display history information for Routing Information Base (RIB) clients, use the **show rib history** command in EXEC mode.

show rib [**afi-all** | **ipv4** | **ipv6**] **history** [**client-id** *client-id*] [**standby**]

Syntax Description

afi-all	(Optional) Specifies all address families.
ipv4	(Optional) Specifies IP Version 4 address prefixes. This is the default.
ipv6	(Optional) Specifies IP Version 6 address prefixes.
client-id *client-id*	(Optional) Specifies the ID of the client. Range for *client-id* argument is 0 to 4294967295.
standby	(Optional) Displays standby information.

Command Default

No default behavior or values

Command Modes

EXEC

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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **show rib history** command to display the list of routes that RIB has sent to various clients.

Task ID

Task ID	Operations
rib	read

Examples

The following is sample output from the **show rib history** command:

```
RP/0/RSP0/CPU0:router# show rib history

JID   Client             Location
229   isis               node0_5_CPU0
  Table ID: 0xe0000000
     S 80.80.80.0/24[1/0]        update, 6 path(s),     04:32:09
     S 100.100.100.0/24[1/0]     update, 1 path(s),     04:32:09
     S 40.40.40.0/24[1/0]        update, 1 path(s),     04:32:09
     S 15.15.15.0/24[1/0]        update, 1 path(s),     04:32:09
JID   Client             Location
260   ospf               node0_5_CPU0
```

```
Table ID: 0xe0000000
   S 80.80.80.0/24[1/0]         update, 6 path(s),     04:32:09
   S 100.100.100.0/24[1/0]      update, 1 path(s),     04:32:09
   S 40.40.40.0/24[1/0]         update, 1 path(s),     04:32:09
   S 15.15.15.0/24[1/0]         update, 1 path(s),     04:32:09
```

This table describes the significant fields shown in the display.

Table 146: show rib history Field Descriptions

Field	Description
JID	Job ID of the client process.
Client	Name of the client process.
Location	Information about where the client process is running.

show rib next-hop

To display registered next-hop notification addresses, use the **show rib next-hop** command in EXEC mode.

show rib [**vrf** {*vrf-name* | **all**}] [**afi-all** | **ipv4** | **ipv6**] [**unicast** | **multicast** | **safi-all**] **next-hop** [*client-name*] [*type interface-path-id* | *ip-address /prefix-length* | *ip-address mask* | **resolved** | **unresolved** | **damped**] [**summary**] [**standby**]

Syntax Description

vrf { *vrf-name* | **all** }	(Optional) Specifies a particular VPN routing and forwarding (VRF) instance or all VRF instances.
afi-all	(Optional) Specifies all address families.
ipv4	(Optional) Specifies IP Version 4 address prefixes. This is the default.
ipv6	(Optional) Specifies IP Version 6 address prefixes.
unicast	(Optional) Specifies unicast address prefixes. This is the default.
multicast	(Optional) Specifies multicast address prefixes.
safi-all	(Optional) Specifies unicast and multicast address prefixes.
client-name	(Optional) Name of the RIB client.
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.
ip-address	(Optional) Network IP address about which routing information should be displayed.
mask	(Optional) Network mask specified in either of two ways: • Network mask can be a four-part, dotted-decimal address. For example, 255.0.0.0 indicates that each bit equal to 1 means the corresponding address bit is a network address. • Network mask can be indicated as a slash (/) and number. For example, /8 indicates that the first 8 bits of the mask are 1s, and the corresponding bits of the address are the network address.
/ *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.
resolved	(Optional) Specifies resolved next-hops.

unresolved	(Optional) Specifies unresolved next-hops.
damped	(Optional) Specifies next-hops that are damped.
summary	(Optional) Specifies a summary of the next-hop information.
standby	(Optional) Displays standby information.

Command Default

No default behavior or values

Command Modes

EXEC

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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **show rib next-hop** command to display the list of next-hops registered by various clients with the RIB and the address and interface through which they are resolved.

Task ID

Task ID	Operations
rib	read

Examples

The following is sample output from the **show rib next-hop** command:

```
RP/0/RSP0/CPU0:router# show rib next-hop

Registered nexthop notifications:

0.0.0.0/0 via 172.29.52.1 - MgmtEth0/RP1/CPU0/0, ospf/node0_RP0_CPU0
172.29.52.1/32 via 172.29.52.1 - MgmtEth0/RP1/CPU0/0, ipv4_static/node0_RP0_CPU0
```

show rib opaques

To display opaque data installed in the Routing Information Base (RIB), use the **show rib opaques** command in EXEC mode.

show rib [**vrf** {*vrf-name* | **all**}] [**afi-all** | **ipv4** | **ipv6**] [**unicast** | **multicast** | **safi-all**] **opaques** {**attribute** | **ip-nexthop** | **ipfrr** | **safi-tunnel** | **summary** | **tunnel-nexthop**} [*rib-client-name*] [**standby**]

Syntax Description

vrf { *vrf-name* | **all** }	(Optional) Specifies a particular VPN routing and forwarding (VRF) instance or all VRF instances.
afi-all	(Optional) Specifies all address families.
ipv4	(Optional) Specifies IP Version 4 address prefixes. This is the default.
ipv6	(Optional) Specifies IP Version 6 address prefixes.
unicast	(Optional) Specifies unicast address prefixes. This is the default.
multicast	(Optional) Specifies multicast address prefixes.
safi-all	(Optional) Specifies unicast and multicast address prefixes.
attribute	Displays opaque attributes installed in the RIB.
ip-nexthop	Displays IP next-hop data installed in the RIB.
ipfrr	Displays IP fast reroute (IPFRR) opaque data installed in the RIB.
safi-tunnel	Displays subaddress family (SAFI) tunnel opaque data installed in the RIB.
summary	Displays a summary of opaque data installed in the RIB.
tunnel-nexthop	Displays tunnel next-hop opaque data installed in the RIB.
rib-client-name	(Optional) Name of the RIB client.
standby	(Optional) Displays standby information.

Command Default

No default behavior or values

Command Modes

EXEC

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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

If information is not used by the RIB server process, it is viewed as opaque data. Use the **show rib opaques** command to display opaque data installed in the RIB.

Task ID

Task ID	Operations
rib	read

Examples

The following is sample output from the **show rib opaques** command:

```
  RP/0/RSP0/CPU0:router# show rib opaques safi-tunnel

  Summary of safi tunnel opaque data in IPv4 RIB:

  Opaque key: 1:10.1.0.2
  Opaque data:
  Tunnel Encap - ifhandle=0x1000180, type=L2TPv3, Params=[Session-id=0x1EB1127C, `
Cookielen=8, Cookie=0xA73A3E0AFCD419A6] Opaque key: 65535:10.0.101.1 Opaque data:


  RP/0/RSP0/CPU0:router# show rib ipv6 opaques tunnel-nexthop

  Summary of 6PE/6VPE IP over tunnel nexthop opaque data in IPv6 RIB:

  Opaque key: 1:::ffff:10.1.0.2
  Opaque key: 65535:::ffff:10.0.101.1
  Opaque key: 65535:::ffff:10.0.101.2
  Opaque key: 65535:::ffff:10.0.101.3
  Opaque key: 65535:::ffff:10.0.101.4
  Opaque key: 65535:::ffff:10.0.101.5
```

This table describes the significant fields shown in the display.

Table 147: show rib opaques Field Descriptions

Field	Description
Opaque key	Unique key for the opaque data as populated by the protocol client.
Opaque data	Data for the given key.

Related Commands

Command	Description
show route, on page 1500	Displays current routes information in the Routing Information Base (RIB).
show ospf routes, on page 1246	Displays Open Shortest Path First (OSPF) topology table.

show rib protocols

To display protocols registered for route addition, use the **show rib protocols** command in EXEC mode.

show rib [**vrf** {*vrf-name* | **all**}] [**afi-all** | **ipv4** | **ipv6**] [**unicast** | **multicast** | **safi-all**] **protocols** [**standby**]

Syntax Description

vrf { *vrf-name* | **all** }	(Optional) Specifies a particular VPN routing and forwarding (VRF) instance or all VRF instances.
afi-all	(Optional) Specifies all address families.
ipv4	(Optional) Specifies IP Version 4 address prefixes. This is the default.
ipv6	(Optional) Specifies IP Version 6 address prefixes.
unicast	(Optional) Specifies unicast address prefixes. This is the default.
multicast	(Optional) Specifies multicast address prefixes.
safi-all	(Optional) Specifies unicast and multicast address prefixes.
standby	(Optional) Displays standby information.

Command Default

If a **vrf** *vrf-name* is not specified, the registered first-hop notification addresses are displayed for the default IPv4 unicast VRF.

Command Modes

EXEC

Command History

Release	Modification
Release 3.7.2	This command was introduced.
Release 3.9.0	Asplain format for 4-byte Autonomous system numbers notation was supported. The input parameters and output were modified to display 4-byte autonomous system numbers and extended communities in either asplain or asdot notations.

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
rib	read

Examples

The following is sample output from the **show rib protocols** command:

```
RP/0/RSP0/CPU0:router# show rib protocols
```

```
Protocol   Handle   Instance
isis         0       rib
connected    1
static       2
local        3
bgp          4       102
ospf         5       1
```

This table describes the significant fields shown in the display.

Table 148: show rib protocols Field Descriptions

Field	Description
Protocol	Name of the protocol.
Handle	Handle assigned to the protocol instance.
Instance	Protocol instance.

show rib recursion-depth-max

To display the maximum recursion depth in the Routing Information Base (RIB), use the **show rib recursion-depth-max** command in EXEC mode.

show rib [**afi-all** | **ipv4** | **ipv6**] **recursion-depth-max** [**standby**]

Syntax Description

afi-all	(Optional) Specifies all address families.
ipv4	(Optional) Specifies IP Version 4 address prefixes. This is the default.
ipv6	(Optional) Specifies IP Version 6 address prefixes.
standby	(Optional) Displays standby information.

Command Default

No default behavior or values

Command Modes

EXEC

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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **show rib recursion-depth-max** command to display the maximum recursion depth for RIB. Recursion depth is the number of next-hops that can be specified.

Task ID

Task ID	Operations
rib	read

Examples

The following is sample output from the **show rib recursion-depth-max** command:

```
RP/0/RSP0/CPU0:router# show rib recursion-depth-max

IPv4:
-----
Maximum recursion depth in RIB:

    Configured: 12
        In Use: 128



IPv6:
-----
Maximum recursion depth in RIB:
```

```
Configured: 12
    In Use: 128
```

This table describes the significant fields shown in the display.

Table 149: show rib recursion-depth-max Field Descriptions

Field	Description
Configured	Value of maximum recursion depth currently configured.
In Use	Value of maximum recursion depth RIB is using. This value can be different from the configured value because RIB has to be restarted after the configuration is changed for the new configuration to be effective.

show rib statistics

To display Routing Information Base (RIB) statistics, use the **show rib statistics** command in EXEC mode.

show rib [**vrf** {*vrf-name* | **all**}] [**afi-all** | **ipv4** | **ipv6**] [**unicast** | **multicast** | **safi-all**] **statistics** [*client-name*] [**standby**]

Syntax Description

vrf { *vrf-name* | **all** }	(Optional) Specifies a particular VPN routing and forwarding (VRF) instance or all VRF instances.
afi-all	(Optional) Specifies all address families.
ipv4	(Optional) Specifies IP Version 4 address prefixes. This is the default.
ipv6	(Optional) Specifies IP Version 6 address prefixes.
unicast	(Optional) Specifies unicast address prefixes. This is the default.
multicast	(Optional) Specifies multicast address prefixes.
safi-all	(Optional) Specifies unicast and multicast address prefixes.
client-name	(Optional) Name of the RIB client.
standby	(Optional) Displays standby information.

Command Default

If **vrf** *vrf-name* is not specified, the registered first-hop notification addresses are displayed for the default IPv4 unicast VRF.

Command Modes

EXEC

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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **show rib statistics** command to display RIB statistics. The statistics include requests sent from the clients to the RIB and the information redistributed to the client.

RIB maintains counters for all requests sent from a client including:

- Route operations
- Table registrations
- Next-hop registrations
- Redistribution registrations
- Attribute registrations
- Synchronization completion

RIB also maintains the results of the requests.

Task ID

Task ID	Operations
rib	read

Examples

The following is sample output from the **show rib statistics** command:

```
RP/0/RSP0/CPU0:router# show rib statistics

RIB Statistics:
 Received 142 batch messages
          137 route operations, 0 attribute operations
          0 opaque operations
          11 complete operations, 0 convergent operations
  Results of the batch message received:
   142 successes
   0 forward references, 0 invalid client id, 0 unknown errors
   0 memory allocation errors, 0 client lookup errors, table lookup errors 0
   0 proto lookup errors, 0 client proto lookup errors
   ipv4_connected/node0_RP0_CPU0 last performed route operation
    with status BATCH_SUCESS at Jun 26 21:43:33.601

 Received 217422 light weight messages
  4 route add requests, 2 route delete requests
  10 protocol registered, 1 protocol unregistered
  0 protocol modify, 0 protocol purged
  14 protocol redistributions, 0 unregistered protocol redistributions
  0 reset protocol redistributions
  3 first hop registered, 1 first hop unregistered
  3 advertisements, 0 unregistered advertisement
  57 bind data, 97 update completes, 217230 other requests
   udp/node0_RP0_CPU0 last performed firsthop lookup operation
    with status success at Jun 27 10:09:59.990

 Received 0 nexthop batch messages
    0 successes
    0 inits
    0 registers, 0 unregisters
    0 register complete, 0 sync unregistered, 0 batch finished
```

This table describes the significant fields shown in the display.

Table 150: show rib statistics Field Descriptions

Field	Description
Received	Statistics received including batch messages and route, attribute, complete, and convergent operations.
Results of the batch message received	Batch message results.
Received *n* light weight messages	Number of lightweight API messages sent from RIB clients.

Field	Description
Received *n* nexthop batch messages	Number of batch API messages sent from RIB clients received by the RIB.

show rib tables

To display all tables known to the Routing Information Base (RIB), use the **show rib tables** command in EXEC mode.

show rib [**afi-all** | **ipv4** | **ipv6**] **tables** [**summary**] [**standby**]

Syntax Description

afi-all	(Optional) Specifies all address families.
ipv4	(Optional) Specifies IP Version 4 address prefixes. This is the default.
ipv6	(Optional) Specifies IP Version 6 address prefixes.
summary	(Optional) Displays summary table information.
standby	(Optional) Displays standby information.

Command Default

No default behavior or values

Command Modes

EXEC

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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **show rib tables** command to display all tables known to the RIB, including table attributes. Attributes include VPN routing and forwarding (VRF) instance, address family, and maximum prefix information.

Task ID

Task ID	Operations
rib	read

Examples

The following is sample output from the **show rib tables** command when entered without an address:

```
RP/0/RSP0/CPU0:router# show rib tables

Codes: N - Prefix Limit Notified, F - Forward Referenced
       D - Table Deleted, C - Table Reached Convergence

VRF             SAFI  Table ID     PrfxLmt   PrfxCnt TblVersion  N F D C
default         uni   0xe0000000   2000000        72        137  N N N Y
default         multi 0xe0100000   2000000         0          0  N N N Y
```

This table describes the significant fields shown in the display.

Table 151: show rib tables Field Descriptions

Field	Description
VRF	Name of the VRF instance.
SAFI	Subaddress family instance.
Table ID	ID of the RIB table.
PrfxLmt	Configured prefix limit for the RIB table.
PrfxCnt	Number of configured prefixes in the RIB table.
TblVersion	Tables version number.
N	Message sent when prefix limit is exceeded.
F	Forward referenced. If Y is indicated, a table has been created by RIB because a client has registered for the table, but RIB has not heard from the router space infrastructure (RSI) about the table. RSI manages the tables.
D	If Y is indicated, the table has been deleted in the RSI but RIB has not cleared the information.
C	Table reached convergence.

show rib trace

To display all Routing Information Base (RIB) library call tracer (ltrace) entries, use the **show rib trace** command in EXEC mode.

show rib [**afi-all** | **ipv4** | **ipv6**] **trace** [**clear** | **counts** | **event-manager** | **startup** | **sync** | **timing**] [**unique** | **wrapping**] [**last** *entries*] [**hexdump**] [**reverse**] [**tailif**] [**stats**] [**verbose**] [**file** *name* **original location** *node-id* | **location** {**all***node-id*}]

Syntax Description

afi-all	(Optional) Specifies all address families.
ipv4	(Optional) Specifies IP Version 4 address prefixes. This is the default.
ipv6	(Optional) Specifies IP Version 6 address prefixes.
counts clear	(Optional) Displays route clear trace entries.
counts	(Optional) Displays counts trace entries.
event-manager	(Optional) Displays RIB event manager trace entries.
startup	(Optional) Displays RIB startup trace entries.
sync	(Optional) Displays client synchronization trace entries.
timing	(Optional) Displays timing trace entries.
unique	(Optional) Displays unique entries with counts.
wrapping	(Optional) Displays wrapping entries.
last *entries*	(Optional) Displays a specified number of the last entries. Range is 1 to 4294967295.
hexdump	(Optional) Displays traces in hexadecimal format.
reverse	(Optional) Displays the latest traces first.
tailif	(Optional) Displays new traces as they are added.
stats	(Optional) Displays statistics.
verbose	(Optional) Displays internal debugging information.
file *name* **original location** *node-id*	(Optional) Displays trace entries for a specific file for the designated node. The *node-id* argument is entered in the *rack/slot/module* notation.
location { **all** | *node-id* }	(Optional) Displays ltrace entries for the designated node. The *node-id* argument is entered in the *rack/slot/module* notation. The **all** keyword displays ltrace entries for all nodes.

Command Default

No default behavior or values

Command Modes

EXEC

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 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
rib	read

Examples

The following is sample output from the **show rib trace** command

```
  RP/0/RSP0/CPU0:router# show rib trace

  1784 wrapping entries (13312 possible, 0 filtered, 1784 total)
  Mar 16 14:59:27.947 rib/ipv4_rib/rib-startup 0/RSP0/CPU0 t1  Create: Management thread
  Mar 16 14:59:27.959 rib/ipv4_rib/rib-startup 0/RSP0/CPU0 t2  Create: Management event
manager
  Mar 16 14:59:28.346 rib/ipv4_rib/rib-io 0/RSP0/CPU0 t1  Initialise: RIB server
  Mar 16 14:59:28.346 rib/ipv4_rib/rib-io 0/RSP0/CPU0 t1  Initialise: Client collection
  Mar 16 14:59:28.676 rib/ipv4_rib/rib-io 0/RSP0/CPU0 t1  Initialise: DB collection
  Mar 16 14:59:28.693 rib/ipv4_rib/rib-io 0/RSP0/CPU0 t1  Initialise: Timer tree
  Mar 16 14:59:28.694 rib/ipv4_rib/rib-io 0/RSP0/CPU0 t1  RUMP: Bind to sysdb
/ipc/gl/ipv4-rib/ for protocol notification
  Mar 16 14:59:29.102 rib/ipv4_rib/rib-startup 0/RSP0/CPU0 t2  Initialise: Debugging routine

  Mar 16 14:59:29.128 rib/ipv4_rib/rib-io 0/RSP0/CPU0 t1  Register: read, select cb functions

  Mar 16 14:59:29.137 rib/ipv4_rib/rib-startup 0/RSP0/CPU0 t1  Register: cerrno DLL name
lib_rib_error.dll
   .
   .
   .
```

show rib vpn-attributes

To display all VPN attributes installed in the Routing Information Base (RIB), use the **show rib vpn-attributes** command in EXEC mode.

show rib [**afi-all** | **ipv4** | **ipv6**] **vpn-attributes** [**summary**] [**standby**]

Syntax Description

afi-all	(Optional) Specifies all address families.
ipv4	(Optional) Specifies IP Version 4 address prefixes.
ipv6	(Optional) Specifies IP Version 6 address prefixes.
summary	(Optional) Displays VPN attribute information.
standby	(Optional) Displays standby information.

Command Default

The default is IPv4 address prefixes.

Command Modes

EXEC

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 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
rib	read

Examples

The following is sample output from the **show rib vpn-attributes** command:

```
RP/0/RSP0/CPU0:router# show rib vpn-attributes

Extended community data in RIB:

Extended community                                Ref count
COST:128:128:41984                                        2
COST:128:129:42240                                        2
COST:128:129:44544                                        1
COST:128:129:169984                                       2
COST:128:129:307200                                       1
EIGRP route-info:0x0:0                                    6
EIGRP route-info:0x8000:0                                 2
EIGRP AD:444:25600                                        2
EIGRP AD:444:25856                                        2
EIGRP AD:444:28160                                        1
```

```
EIGRP AD:444:51200                                      1
EIGRP AD:444:153600                                     2
EIGRP RHB:255:0:16384                                   2
EIGRP RHB:255:1:16384                                   5
EIGRP RHB:255:1:256000                                  1
EIGRP LM:0x0:1:1500                                     3
EIGRP LM:0x0:1:1514                                     2
EIGRP LM:0x0:1:4470                                     3
EIGRP AR:0:192.168.0.13                                 6
EIGRP PM:11:0                                           6


MVPN attribute data in RIB:

MVPN Attribute                                      Ref count
0:0:1:f4:0:0:0:1:1:1:1:1                                    1
0:0:2:bc:0:0:0:1:3:3:3:3                                   10
0:0:2:bc:0:0:0:1:3:3:3:4                                    2
```

This table describes the significant fields shown in the display.

Table 152: show rib vpn-attributes Field Descriptions

Field	Description
Extended Community	Extended community added by the protocol clients.
Ref Count	Number of routes referring to the same extended community.
MVPN Attribute	Connector attribute added by BGP to support MVPNs.
Ref Count	Number of routes referring to the same extended community.

show rib vrf

To display all VRF table information in the Routing Information Base (RIB), use the **show rib vrf** command in EXEC mode.

show rib vrf {*vrf-name* | **all**} [**ipv4**] [**ipv6**] [**afi-all**] [**firsthop**] [**next-hop**] [**opaques**] [**protocols**] [**statistics** *name*]

Syntax Description

vrf { *vrf-name* | **all** }	(Optional) Specifies a particular VPN routing and forwarding (VRF) instance or all VRF instances.
ipv4	(Optional) Specifies IP Version 4 address prefixes.
ipv6	(Optional) Specifies IP Version 6 address prefixes.
afi-all	(Optional) Specifies all address families.
firsthop	(Optional) Specifies registered first-hop notification addresses
next-hop	(Optional) Specifies registered next-hop notification addresses.
opaques	(Optional) Specifies opaque data installed in the RIB.
protocols	(Optional) Specifies registered protocols.
statistics *name*	(Optional) Specifies RIB statistics for the given name.

Command Default

No default behavior or values

Command Modes

EXEC

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 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
ipv4	read

Examples

The following example shows output from the **show rib vrf all statistics** command:

```
RP/0/RSP0/CPU0:router# show rib vrf all statistics
RP/0/RSP0/CPU0:router#
```

Related Commands

Command	Description
show rib, on page 1466	Displays RIB information.

show route

To display the current routes in the Routing Information Base (RIB), use the **show route** command in EXEC mode.

show route [**vrf** {*vrf-name* | **all**}] [**afi-all** | **ipv4** | **ipv6**] [**unicast** | **multicast** | **safi-all**] [**protocol** [*instance*] | *ip-address* [*mask*] | *ip-address*/*prefix-length*] [**standby**] [**detail**]

Syntax Description

vrf { *vrf-name* | **all** }	(Optional) Specifies a particular VPN routing and forwarding (VRF) instance or all VRF instances.
afi-all	(Optional) Specifies all address families.
ipv4	(Optional) Specifies IP Version 4 address prefixes. This is the default.
ipv6	(Optional) Specifies IP Version 6 address prefixes.
unicast	(Optional) Specifies unicast address prefixes. This is the default.
multicast	(Optional) Specifies multicast address prefixes.
safi-all	(Optional) Specifies unicastand multicast address prefixes.
protocol	(Optional) Name of a routing protocol. If you specify a routing protocol, use one of the following keywords: • **bgp** • **eigrp** • **isis** • **ospf** • **rip** • **static** • **local** • **connected**
instance	(Optional) Number or name used to identify an instance of the specified protocol.
ip-address	(Optional) Network IP address about which routing information should be displayed.
mask	(Optional) Network mask specified in either of two ways: • Network mask can be a four-part, dotted-decimal address. For example, 255.0.0.0 indicates that each bit equal to 1 means the corresponding address bit is a network address. • Network mask can be indicated as a slash (/) and number. For example, /8 indicates that the first 8 bits of the mask are 1s, and the corresponding bits of the address are the network address.
/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.

standby	(Optional) Displays standby information.
detail	(Optional) Displays detailed information for the specified prefix.

Command Default

If a **vrf** *vrf-name* is not specified, routes are displayed for the default IPv4 unicast VRF.

Command Modes

EXEC

Command History

Release	Modification
Release 3.7.2	This command was introduced.
Release 3.9.0	Asplain format for 4-byte Autonomous system numbers notation was supported. The input parameters and output were modified to display 4-byte autonomous system numbers and extended communities in either asplain or asdot notations.

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.

When the **afi-all** keyword is used, the *ip-address* and *mask* arguments are not available.

The **topology** keyword must be accompanied by the **ipv4 multicast** keywords, except when the **afi-all** keyword or the **safi-all** keyword is specified.

Task ID

Task ID	Operations
rib	read

Examples

The following is sample output from the **show route** command when entered without an address:

```
RP/0/RSP0/CPU0:router# show route

Codes: C - connected, S - static, R - RIP, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP
       i - ISIS, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, su - IS-IS summary null, * - candidate default
       U - per-user static route, o - ODR, L - local, G  - DAGR
       A - access/subscriber, (!) - FRR Backup path

Gateway of last resort is 1.0.0.1 to network 0.0.0.0

S*   0.0.0.0/0 [1/0] via 1.0.0.1, 13:14:59
C    1.0.0.0/16 is directly connected, 13:14:59, MgmtEth0/5/CPU0/0
L    1.0.14.15/32 is directly connected, 13:14:59, MgmtEth0/5/CPU0/0
C    3.2.3.0/24 is directly connected, 00:04:39, GigabitEthernet0/3/0/0
L    3.2.3.2/32 is directly connected, 00:04:39, GigabitEthernet0/3/0/0
O E2 5.2.5.0/24 [110/20] via 3.3.3.1, 00:04:20, GigabitEthernet0/3/0/0
O E2 6.2.6.0/24 [110/20] via 3.3.3.1, 00:04:20, GigabitEthernet0/3/0/0
C    7.2.7.0/24 is directly connected, 00:04:20, GigabitEthernet0/3/0/7
L    7.2.7.2/32 is directly connected, 00:04:20, GigabitEthernet0/3/0/7
```

```
O E2 8.2.8.0/24 [110/20] via 3.3.3.1, 00:04:20, GigabitEthernet0/3/0/0

C    10.3.0.0/16 is directly connected, 13:14:59, GigabitEthernet0/0/0/0
L    10.3.0.2/32 is directly connected, 13:14:59, GigabitEthernet0/0/0/0
```

This table describes the significant fields shown in the display.

Table 153: show route Field Descriptions

Field	Description
S*	Code indicating how the route was derived. See the code legend preceding the output. In this case, the route was derived from a static (candidate default).
[1/0]	First number in the brackets is the administrative distance of the information source; the second number is the metric for the route.
1.0.0.0/16	Address and prefix length of the remote network.
MgmtEthernet 0/5/CPU0/0	Specifies the interface through which the specified network can be reached.
C	Code indicating how the route was derived. See the code legend preceding the output. In this case, the route was connected.
L	Code indicating how the route was derived. See the code legend preceding the output. In this case, the route was local.
O	Code indicating how the route was derived. See the code legend preceding the output. In this case, the route was on-demand routing (ODR).
E2	Code indicating how the route was derived. See the code legend preceding the output. In this case, the route was OSPF external type 2.
8.2.8.0/24	Address and prefix length of the remote network connected to the static route.
via 3.3.3.1	Specifies the address of the next router to the remote network.
13:14:59	Specifies the last time the route was updated.

When you specify that you want information about a particular network, more detailed statistics are displayed. The following is sample output from the **show route** command when entered with an IP address:

```
RP/0/RSP0/CPU0:router# show route 10.0.0.0

 Routing entry for 10.0.0.0/16
   Known via "connected", distance 0, metric 0 (connected)
   Installed Mar 22 22:10:20.906
   Routing Descriptor Blocks
     directly connected, via GigabitEthernet0/0/0/0
       Route metric is 0
   No advertising protos.
```

Intermediate System-to-Intermediate System (IS-IS) includes an IP address typed length value (TLV) in its link-state packet (LSP) that helps identify the node injecting the route into the network. The IS-IS node uses one of its own interface addresses in this TLV. A loopback address is preferred among interfaces configured under IS-IS. When other networking devices calculate IP routes, they can store the IP address as the originator address with each route in the routing table.

The following example shows the output from the **show route** command for a specific IP address on a router configured with IS-IS. Each path that is shown under the Routing Descriptor Blocks report displays two IP addresses. The first address (10.0.0.9) is the next-hop address; the second is the originator IP address from the advertising IS-IS router.

```
RP/0/RSP0/CPU0:router# show route 10.0.0.1

 Routing entry for 10.0.0.0/8
 Known via "isis", distance 115, metric 10, type level-2
      Installed Jan 22 09:26:56.210
      Routing Descriptor Blocks:
      * 10.0.0.9, from 10.0.0.9, via GigabitEthernet2/1
       Route metric is 10
    No advertising protos.
```

This table describes the significant fields shown in the display.

Table 154: show route with IP Address Field Descriptions

Field	Description
Routing entry for	Network address and mask.
Known via	Indicates how the route was derived.
distance	Administrative distance of the information source.
metric	Route value assigned by the routing protocol.
type	IS-IS type level.
Routing Descriptor Blocks:	Displays the next-hop IP address followed by the information source.
from ... via ...	First address is the next-hop IP address, and the other is the information source. This report is followed by the interface for this route.
Route metric	Best metric for this Routing Descriptor Block.

No advertising protos.	Indicates that no other protocols are advertising the route to their redistribution consumers. If the route is being advertised, protocols are listed in the following manner:
	
	`Redist Advertisers:`
	`  isis p`
	`  ospf 43`

The following example illustrates the **show route** command with the **topology** *topo-name* keyword and argument specified:

```
RP/0/RSP0/CPU0:router# show route ipv4 multicast topology green

  Codes: C - connected, S - static, R - RIP, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP
       i - ISIS, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, su - IS-IS summary null, * - candidate default
       U - per-user static route, o - ODR, L - local, G  - DAGR
       A - access/subscriber, (!) - FRR Backup path

  Gateway of last resort is not set

  i L1 10.1.102.0/24 [115/20] via 10.1.102.41, 1w4d, GigabitEthernet0/1/0/0.1
  i L1 10.3.3.0/24 [115/20] via 10.1.102.41, 1w4d, GigabitEthernet0/1/0/0.1
  i L1 192.168.0.40/32 [115/20] via 10.1.102.41, 1w4d, GigabitEthernet0/1/0/0.1
```

This example is a sample **show route detail** command output that displays path ID and backup-path ID information:

```
RP/0/RSP0/CPU0:router#show route 10.1.1.3 detail
Routing entry for 10.1.1.16/32
  Known via "ospf 2", distance 110, metric 21, type intra area
  Installed Oct 28 16:07:05.752 for 00:01:56
  Routing Descriptor Blocks
    40.1.10.1, from 10.1.1.16, via Bundle-Ether10, Protected
      Route metric is 21
      Label: None
      Tunnel ID: None
      Extended communities count: 0
      Path id:2       Path ref count:0
      Backup path id:33
    200.40.1.101, from 10.1.1.16, via Bundle-Ether1.1, Protected
      Route metric is 21
      Label: None
      Tunnel ID: None
      Extended communities count: 0
      Path id:1       Path ref count:0
      Backup path id:33
    100.100.2.1, from 10.1.1.16, via TenGigE0/2/0/3.1, Backup
      Route metric is 0
      Label: None
      Tunnel ID: None
      Extended communities count: 0
      Path id:33       Path ref count:2
  Route version is 0xe (14)
  No local label
  IP Precedence: Not Set
  QoS Group ID: Not Set
  Route Priority: RIB_PRIORITY_NON_RECURSIVE_LOW (6) SVD Type RIB_SVD_TYPE_LOCAL
  No advertising protos.
```

Related Commands

Command	Description
show interfaces	Lists interface information.
show route summary, on page 1525	Displays the current contents of the routing table in summary format.

Command	Description
show rib opaques, on page 1483	Displays opaque data installed in the Routing Information Base (RIB).
show ospf routes, on page 1246	Displays Open Shortest Path First (OSPF) topology table.

show route backup

To display backup routes from the Routing Information Base (RIB), use the **show route backup** command in EXEC mode.

show route [**vrf** {*vrf-name* | **all**}] [**afi-all** | **ipv4** | **ipv6**] [**unicast** | **multicast** | {**topology** *topo-name*} | **safi-all**] **backup** [*ip-address [mask] ip-address /prefix-length]*][**standby**]

Syntax Description

vrf { *vrf-name* | **all** }	(Optional) Specifies a particular VPN routing and forwarding (VRF) instance or all VRF instances.
afi-all	(Optional) Specifies all address families.
ipv4	(Optional) Specifies IP Version 4 address prefixes.
ipv6	(Optional) Specifies IP Version 6 address prefixes.
unicast	(Optional) Specifies unicast address prefixes.
multicast	(Optional) Specifies multicast address prefixes.
safi-allsafi-all	(Optional) Specifies unicast and multicast address prefixes.
ip-address	(Optional) Network IP address about which backup routing information should be displayed.
mask	(Optional) Network mask specified in either of two ways: • Network mask can be a four-part, dotted decimal address. For example, 255.0.0.0 indicates that each bit equal to 1 means the corresponding address bit is a network address. • Network mask can be indicated as a slash (/) and number. For example, /8 indicates that the first 8 bits of the mask are ones, and the corresponding bits of the address are the network address.
/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.
standby	(Optional) Displays standby information.

Command Default

If a **vrf** *vrf-name* is not specified, the backup routes from the RIB are displayed for the default IPv4 unicast VRF.

Command Modes

EXEC

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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **show route backup** command to display information about routes that have been installed into the RIB as backup routes. This command also displays information about the currently selected active route for which there is a backup.

When the **afi-all** keyword is used, the *ip-address* and *mask* arguments are not available.

The **topology** keyword must be accompanied by the **ipv4 multicast** keywords, except when the **afi-all** keyword or the **safi-all** keyword is specified.

Task ID

Task ID	Operations
rib	read

Examples

The following is sample output from the **show route backup** command:

```
RP/0/RSP0/CPU0:router# show route backup

Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP
       i - ISIS, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, su - IS-IS summary null, * - candidate default
       U - per-user static route, o - ODR, L - local
S    172.73.51.0/24 is directly connected, 2d20h, GigabitEthernet 4/0/0/1
                Backup  O E2 [110/1] via 10.12.12.2, GigabitEthernet 3/0/0/1
```

This table describes the significant fields shown in the display.

Table 155: show route backup Field Descriptions

Field	Description
S	Code indicating how the route was derived. See the legend of the codes preceding the output.
172.73.51.0/24	IP address and length of the route.
2d20h	Time (in hh:mm:ss) since the route was installed in the RIB.
GigabitEthernet4/0/0/1	Outbound interface for the route.
Backup	Identifies the entry as a backup version of the route, typically installed by a different routing protocol.
O	Code indicating how the route was derived. See the code legend preceding the output.

Field	Description
E2	Code for the type of route. This code is relevant only for OSPF and IS-IS routes. The codes for an OSPF route can be: none—intra-area route IA—interarea route E1—external type 1 E2—external type 2 N1—NSSA external type 1 N2—NSSA external type 2 The codes for an IS-IS route can be: L1—level 1 L2—level 2 ia—interarea su—summary route
[110/1]	Distance and metric for the route.
10.12.12.2	IP address of next-hop on the route.
GigabitEthernet3/0/0/1	Outbound interface for the OSPF version of the route.

Related Commands

Command	Description
show route, on page 1500	Displays the current routes in the RIB.

show route best-local

To display the best local address to use for return packets from the given destination, use the **show route best-local** command in EXEC mode.

show route [**vrf** {*vrf-name* | **all**}] [**ipv4** | **ipv6**] [**unicast** | **multicast** | {**topology** *topo-name*} | **safi-all**] **best-local** *ip-address* [**standby**]

Syntax Description

vrf { *vrf-name* | **all** }	(Optional) Specifies a particular VPN routing and forwarding (VRF) instance or all VRF instances.
ipv4	(Optional) Specifies IP Version 4 address prefixes.
ipv6	(Optional) Specifies IP Version 6 address prefixes.
unicast	(Optional) Specifies unicast address prefixes.
multicast	(Optional) Specifies multicast address prefixes.
safi-all	(Optional) Specifies unicast and multicast address prefixes.
ip-address	IP address about which best local information should be displayed.
standby	(Optional) Displays standby information.

Command Default

If a **vrf** *vrf-name* is not specified, the best local address is displayed for the default IPv4 unicast VRF.

Command Modes

EXEC

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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **show route best-local** command to display information about the best local routes in the routing table.

The **topology** keyword must be accompanied by the **ipv4 multicast** keywords, except when the **afi-all** keyword or the **safi-all** keyword is specified.

Task ID

Task ID	Operations
rib	read

Examples

The following is sample output from the **show route best-local** command:

```
RP/0/RSP0/CPU0:router# show route best-local 10.12.12.1/32

Routing entry for 10.12.12.1/32
  Known via "local", distance 0, metric 0 (connected)
  Routing Descriptor Blocks
    10.12.12.1 directly connected, via GigabitEthernet3/0/0/1
      Route metric is 0
```

This table describes the significant fields shown in the display.

Table 156: show route best-local Field Descriptions

Field	Description
Routing entry for	Identifies the requested IP address.
Known via	Indicates how the route was derived.
distance	Administrative distance of the information source.
metric	Route value assigned by the routing protocol.
Routing Descriptor Blocks:	Displays the next-hop IP address followed by the information source.
10.12.12.1 Directly connected ... via ...	First address is the next-hop IP address, followed by a report that it is directly connected. This report is followed by the interface for this route.

Related Commands

Command	Description
show route local, on page 1513	Displays local addresses installed in the RIB as a receive entry.

show route connected

To display the current connected routes of the routing table, use the **show route connected** command in EXEC mode.

show route [**vrf** {*vrf-name* | **all**}] [**afi-all** | **ipv4** | **ipv6**] [**unicast** | **multicast** | {**topology** *topo-name*} | **safi-all**] **connected** [**standby**]

Syntax Description

vrf { *vrf-name* | **all** }	(Optional) Specifies a particular VPN routing and forwarding (VRF) instance or all VRF instances.
afi-all	(Optional) Specifies all address families.
ipv4	(Optional) Specifies IP Version 4 address prefixes.
ipv6	(Optional) Specifies IP Version 6 address prefixes.
unicast	(Optional) Specifies unicast address prefixes.
multicast	(Optional) Specifies multicast address prefixes.
safi-all	(Optional) Specifies unicast and multicast address prefixes.
standby	(Optional) Displays standby information.

Command Default

If a **vrf** *vrf-name* is not specified, the current connected routes of the routing table are displayed for the default IPv4 unicast VRF.

The **topology** keyword must be accompanied by the **ipv4 multicast** keywords, except when the **afi-all** keyword or the **safi-all** keyword is specified.

Command Modes

EXEC

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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **show route connected** command to display information about connected routes in the routing table.

Task ID

Task ID	Operations
rib	read

Examples

The following is sample output from the **show route connected** command:

```
RP/0/RSP0/CPU0:router# show route connected

C    1.68.0.0/16 is directly connected, 13:43:40, MgmtEth0/5/CPU0/0
C    3.3.3.0/24 is directly connected, 00:23:23, GigabitEthernet0/3/0/0
C    7.7.7.0/24 is directly connected, 00:33:00, GigabitEthernet0/3/0/7
C    10.0.0.0/16 is directly connected, 13:43:40, GigabitEthernet0/0/0/0
C    10.10.10.0/30 is directly connected, 13:43:40, Loopback0
C    11.11.11.0/24 is directly connected, 13:43:40, Loopback11
```

This table describes the significant fields shown in the display.

Table 157: show route connected Field Descriptions

Field	Description
C	Code to indicate the route is connected.
1.68.0.0/16	IP address and length of the route.
13:43:40	Time (in hh:mm:ss) since the route was installed in the RIB.
MgmtEth0/5/CPU0/0	Outbound interface for the route.

Related Commands

Command	Description
show route summary, on page 1525	Displays the current contents of the RIB.

show route local

To display local routes receiving routing updates from the Routing Information Base (RIB), use the **show route local** command in EXEC mode.

show route [**vrf** {*vrf-name* | **all**}] [**afi-all** | **ipv4** | **ipv6**] [**unicast** | **multicast** | {**topology** *topo-name*} | **safi-all**] **local** [*type interface -path-id*] [**standby**]

Syntax Description

vrf { *vrf-name* | **all** }	(Optional) Specifies a particular VPN routing and forwarding (VRF) instance or all VRF instances.
afi-all	(Optional) Specifies all address families.
ipv4	(Optional) Specifies IP Version 4 address prefixes.
ipv6	(Optional) Specifies IP Version 6 address prefixes.
unicast	(Optional) Specifies unicast address prefixes.
multicast	(Optional) Specifies multicast address prefixes.
safi-all	(Optional) Specifies unicast and multicast address prefixes.
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.
standby	(Optional) Displays standby information.

Command Default

If a **vrf** *vrf-name* is not specified, the local routes receiving updates from the RIB are displayed for the default IPv4 unicast VRF.

Command Modes

EXEC

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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **show route local** command to display information about local routes in the routing table.

The **topology** keyword must be accompanied by the **ipv4 multicast** keywords, except when the **afi-all** keyword or the **safi-all** keyword is specified.

Task ID

Task ID	Operations
rib	read

Examples

The following is sample output from the **show route local** command:

```
RP/0/RSP0/CPU0:router# show route local

L    10.10.10.1/32 is directly connected, 00:14:36, Loopback0
L    10.91.36.98/32 is directly connected, 00:14:32, GigabitEthernet6/0/0/1
L    172.22.12.1/32 is directly connected, 00:13:35, GigabitEthernet3/0/0/1
L    192.168.20.2/32 is directly connected, 00:13:27, GigabitEthernet4/0/0/1
L    10.254.254.1/32 is directly connected, 00:13:26, GigabitEthernet5/0/0/1
```

This table describes the significant fields shown in the display.

Table 158: show route local Field Descriptions

Field	Description
L	Code to indicate the route is local.
10.10.10.1/32	IP address and length of the route.
00:14:36	Time (in hh:mm:ss) since the route was installed in the RIB.
Loopback0	Outbound interface for the route.

Related Commands

Command	Description
show route connected, on page 1511	Displays information about all clients that have registered with the RIB as protocols.

show route longer-prefixes

To display the current routes in the Routing Information Base (RIB) that share a given number of bits with a given network, use the **show route longer-prefixes** command in EXEC mode.

show route [**vrf** {*vrf-name* | **all**}] [**ipv4** | **ipv6**] [**unicast** | **multicast** | {**topology** *topo-name*} | **safi-all**] **longer-prefixes** {*ip-address mask ip-address/prefix-length*} [**standby**]

Syntax Description

vrf { *vrf-name* | **all** }	(Optional) Specifies a particular VPN routing and forwarding (VRF) instance or all VRF instances.
ipv4	(Optional) Specifies IP Version 4 address prefixes.
ipv6	(Optional) Specifies IP Version 6 address prefixes.
unicast	(Optional) Specifies unicast address prefixes.
multicast	(Optional) Specifies multicast address prefixes.
safi-all	(Optional) Specifies unicast and multicast address prefixes.
ip-address	Network IP address about which routing information should be displayed.
mask	Network mask specified in either of two ways: • Network mask can be a four-part, dotted-decimal address. For example, 255.0.0.0 indicates that each bit equal to 1 means the corresponding address bit is a network address. • Network mask can be indicated as a slash (/) and number. For example, /8 indicates that the first 8 bits of the mask are 1s, and the corresponding bits of the address are the network address.
/ *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.
standby	(Optional) Displays standby information.

Command Default

If a **vrf** *vrf-name* is not specified, the current routes in the RIB sharing a specified number of bits with a network are displayed for the default IPv4 unicast VRF.

Command Modes

EXEC

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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **show route longer-prefixes** command to troubleshoot forwarding problems whose cause may be a long prefix.

The **topology** keyword must be accompanied by the **ipv4 multicast** keywords, except when the **afi-all** keyword or the **safi-all** keyword is specified.

Task ID

Task ID	Operations
rib	read

Examples

The following is sample output from the **show route longer-prefixes** command:

```
RP/0/RSP0/CPU0:router# show route longer-prefixes 172.16.0.0/8

Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP
       i - ISIS, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, su - IS-IS summary null, * - candidate default
       U - per-user static route, o - ODR, L - local

L    172.29.52.70/32 is directly connected, 4d15h, MgmtEth0/RSP0/CPU0/0
L    172.29.52.71/32 is directly connected, 4d15h, MgmtEth0/RP1/CPU0/0
L    172.29.52.72/32 [0/0] via 172.29.52.72, 4d15h, MgmtEth0/RSP0/CPU0/0
```

This table describes the significant fields shown in the display.

Table 159: show route longer-prefixes Field Descriptions

Field	Description
172.29.52.70/32	IP address and length of the route.
4d15h	Time (in hh:mm:ss or *n*d*n*h) since the route was installed in the RIB.
MgmtEth0/RSP0 /CPU0/0	Outbound interface for the route.

Related Commands

Command	Description
router static	Establishes a static route.
show interfaces	Lists interface information.
show route summary, on page 1525	Displays the current contents of the routing table in summary format.

show route next-hop

To filter routes by the next-hop address or interface, use the **show route next-hop** command in EXEC mode.

show route [**vrf** {*vrf-name* | **all**}] [**ipv4** | **ipv6**] [**unicast** | **multicast** | {**topology** *topo-name*} | **safi-all**] **next-hop** [*ip-address*][[**standby**]]

Syntax Description

vrf { *vrf-name* | **all** }	(Optional) Specifies a particular VPN routing and forwarding (VRF) instance or all VRF instances.
ipv4	(Optional) Specifies IP Version 4 address prefixes.
ipv6	(Optional) Specifies IP Version 6 address prefixes.
unicast	(Optional) Specifies unicast address prefixes.
multicast	(Optional) Specifies multicast address prefixes.
safi-all	(Optional) Specifies unicast and multicast address prefixes.
ip-address	(Optional) IP address about which next-hop information is to be displayed.
standby	(Optional) Displays standby information.

Command Default

If a **vrf** *vrf-name* is not specified, the next-hop gateway or host is displayed for the default IPv4 unicast VRF.

Command Modes

EXEC

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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **show route next-hop** command to find all routes going through a next-hop address or interface.

The **topology** keyword must be accompanied by the **ipv4 multicast** keywords, except when the **afi-all** keyword or the **safi-all** keyword is specified.

Task ID

Task ID	Operations
rib	read

Examples

The following is sample output from the **show route next-hop** command filtering routes on the next-hop address:

```
RP/0/RSP0/CPU0:router# show route next-hop  1.68.0.1

Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP
i - ISIS, L1 - IS-IS level-1, L2 - IS-IS level-2
ia - IS-IS inter area, su - IS-IS summary null, * - candidate default
U - per-user static route, o - ODR, L - local

Gateway of last resort is 1.68.0.1 to network 0.0.0.0

S* 0.0.0.0/0 [1/0] via 1.68.0.1, 15:01:49
S 223.255.254.254/32 [1/0] via 1.68.0.1, 15:01:49
```

The following is sample output from the **show route next-hop** command filtering routes on the next-hop interface:

```
RP/0/RSP0/CPU0:router# show route next-hop GigabitEthernet 0/1/0/2

Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP
i - ISIS, L1 - IS-IS level-1, L2 - IS-IS level-2
ia - IS-IS inter area, su - IS-IS summary null, * - candidate default
U - per-user static route, o - ODR, L - local

Gateway of last resort is 1.68.0.1 to network 0.0.0.0

C 11.1.1.0/24 is directly connected, 15:01:46, GigabitEthernet0/1/0/2
L 11.1.1.2/32 is directly connected, 15:01:46, GigabitEthernet0/1/0/2
```

This table describes the significant fields shown in the display.

Table 160: show route next-hop Field Descriptions

Field	Description
11.1.1.0/24	IP address and length of the route.
15:01:46	Time (in hh:mm:ss or *n*d*n*h) since the route was installed in the RIB.
GigabitEthernet0/1/0/2	Outbound interface for the route.

Related Commands

Command	Description
show route, on page 1500	Displays the current contents of the routing table.

show route quarantined

To display mutually recursive (looping) routes, use the **show route quarantined** command in EXEC mode.

show route [**vrf** {*vrf-name* | **all**}] [**ipv4** | **ipv6**] [**unicast** | **multicast** | {} | **safi-all**] **quarantined** [*ip-address/prefix-length]ip-address mask*] [**standby**]

Syntax Description

vrf { *vrf-name* | **all** }	(Optional) Specifies a particular VPN routing and forwarding (VRF) instance or all VRF instances.
ipv4	(Optional) Specifies IP Version 4 address prefixes.
ipv6	(Optional) Specifies IP Version 6 address prefixes.
unicast	(Optional) Specifies unicast address prefixes.
multicast	(Optional) Specifies multicast address prefixes.
safi-all	(Optional) Specifies unicast and multicast address prefixes.
ip-address	(Optional) IP address about which looping routes information is to be displayed.
/ *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.
ip-address mask	(Optional) Network mask applied to the *ip-address* argument.
standby	(Optional) Displays standby information.

Command Default

If a **vrf** *vrf-name* is not specified, the next-hop gateway or host is displayed for the default IPv4 unicast VRF.

Command Modes

EXEC

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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

RIB quarantining detects mutually recursive routes and quarantines the last route that actually completes the mutual recursion. The quarantined route is periodically evaluated to see if the mutual recursion has gone away. If the recursion still exists, the route remains quarantined. If the recursion has gone away, the route is released from quarantine.

Use the **show route quarantined** command to display mutually recursive (looping) routes.

The **topology** keyword must be accompanied by the **ipv4 multicast** keywords, except when the **afi-all** keyword or the **safi-all** keyword is specified.

Task ID

Task ID	Operations
rib	read

Examples

The following is sample output from the **show route quarantined** command:

```
RP/0/RSP0/CPU0:routerr# show route quarantined

Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP
       i - ISIS, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, su - IS-IS summary null, * - candidate default
       U - per-user static route, o - ODR, L - local

S    10.10.109.1/32 [1/0] via 10.10.34.1, 00:00:01 (quarantined)
                    [1/0] via 10.10.37.1, 00:00:01 (quarantined)
                    [1/0] via 10.10.60.1, 00:00:01 (quarantined)
                    [1/0] via 10.10.68.1, 00:00:01 (quarantined)
                    [1/0] via 10.10.91.1, 00:00:01 (quarantined)
                    [1/0] via 10.10.93.1, 00:00:01 (quarantined)
                    [1/0] via 10.10.97.1, 00:00:01 (quarantined)
S    10.0.0.0/8 [1/0] via 11.11.11.11, 00:01:29 (quarantined)
S    10.10.0.0/16 [1/0] via 11.11.11.11, 00:01:29 (quarantined)
S    10.10.10.0/24 [1/0] via 11.11.11.11, 00:01:29 (quarantined)
S    10.10.10.10/32 [1/0] via 11.11.11.11, 00:00:09 (quarantined)
```

This table describes the significant fields shown in the display.

Table 161: show route quarantined Field Descriptions

Field	Description
10.10.109.1/32	IP address and length of the route.
[1/0]	Distance and metric for the route.
via 10.10.34.1	IP address of next-hop on the route.
00:00:01	Time (in hh:mm:ss or *n*d*n*h) since the route was installed in the RIB.
(quarantined)	Shows that the route is quarantined.

Related Commands

Command	Description
show route, on page 1500	Displays the current contents of the routing table.

show route resolving-next-hop

To display the next-hop gateway or host to a destination address, use the **show route resolving-next-hop** command in EXEC mode.

show route [**vrf** {*vrf-name* | **all**}] [**ipv4** | **ipv6**] [**unicast** | **multicast** | {**topology** *topo-name*} | **safi-all**] **resolving-next-hop** *ip-address* [**standby**]

Syntax Description

vrf { *vrf-name* | **all** }	(Optional) Specifies a particular VPN routing and forwarding (VRF) instance or all VRF instances.
ipv4	(Optional) Specifies IP Version 4 address prefixes.
ipv6	(Optional) Specifies IP Version 6 address prefixes.
unicast	(Optional) Specifies unicast address prefixes.
multicast	(Optional) Specifies multicast address prefixes.
safi-all	(Optional) Specifies unicast and multicast address prefixes.
ip-address	IP address about which resolved next-hop information is to be displayed.
standby	(Optional) Displays standby information.

Command Default

If a **vrf** *vrf-name* is not specified, the next-hop gateway or host is displayed for the default IPv4 unicast VRF.

Command Modes

EXEC

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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **show route resolving-next-hop** command to perform a recursive route lookup on the supplied destination address and return information on the next immediate router (next hop) to the destination.

The **topology** keyword must be accompanied by the **ipv4 multicast** keywords, except when the **afi-all** keyword or the **safi-all** keyword is specified.

Task ID

Task ID	Operations
rib	read

Examples

The following is sample output from the **show route resolving-next-hop** command:

```
RP/0/RSP0/CPU0:router# show route resolving-next-hop 10.1.1.1

Nexthop matches 10.1.1.1/32
  Known via "local", distance 0, metric 0 (connected)
  Installed Aug 22 01:57:08.514
  Directly connected nexthops
    10.1.1.1 directly connected, via Loopback0
      Route metric is 0
```

This table describes the significant fields shown in the display.

Table 162: show route resolving-next-hop Field Descriptions

Field	Description
Known via	Name of the routing protocol that installed the matching route.
Route metric is	Metric of the route.

Related Commands

Command	Description
show route, on page 1500	Displays the current contents of the routing table.

show route static

To display the current static routes of the Routing Information Base (RIB), use the **show route static** command in EXEC mode.

show route [**vrf** {*vrf-name* | **all**}] [**afi-all** | **ipv4** | **ipv6**] [**unicast** | **multicast** | {**topology** *topo-name*} | **safi-all**] **static** [**standby**]

Syntax Description

vrf { *vrf-name* | **all** }	(Optional) Specifies a particular VPN routing and forwarding (VRF) instance or all VRF instances.
afi-all	(Optional) Specifies all address families.
ipv4	(Optional) Specifies IP Version 4 address prefixes.
ipv6	(Optional) Specifies IP Version 6 address prefixes.
unicast	(Optional) Specifies unicast address prefixes.
multicast	(Optional) Specifies multicast address prefixes.
safi-all	(Optional) Specifies unicast and multicast address prefixes.
standby	(Optional) Displays standby information.

Command Default

If a **vrf** *vrf-name* is not specified, the current static routes of the RIB are displayed for the default IPv4 unicast VRF.

Command Modes

EXEC

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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **show route static** command to display information about static routes in the routing table.

The **topology** keyword must be accompanied by the **ipv4 multicast** keywords, except when the **afi-all** keyword or the **safi-all** keyword is specified.

Task ID

Task ID	Operations
rib	read

Examples

The following is sample output from the **show route static** command:

```
RP/0/RSP0/CPU0:router# show route static

S    10.1.1.0/24 is directly connected, 00:54:05, GigabitEthernet3/0/0/1
S    192.168.99.99/32 [1/0] via 10.12.12.2, 00:54:04
```

This table describes the significant fields shown in the display.

Table 163: show route static Field Descriptions

Field	Description
S	Code to indicate the route is static.
10.1.1.0/24	IP address and distance for the route.
00:54:05	Time (in hh:mm:ss) since the route was installed in the RIB.
GigabitEthernet3/0/0/1	Outbound interface for the route.
[1/0]	Distance and metric for the route.

Related Commands

Command	Description
show route, on page 1500	Displays the current contents of the routing table.

show route summary

To display the current contents of the Routing Information Base (RIB), use the **show route summary** command in EXEC mode.

show route [**vrf** {*vrf-name* | **all**}] [**afi-all** | **ipv4** | **ipv6**] [**unicast** | **multicast** | {**topology** *topo-name*} | **safi-all**] **summary** [**detail**] [**standby**]

Syntax Description

vrf { *vrf-name* | **all** }	(Optional) Specifies a particular VPN routing and forwarding (VRF) instance or all VRF instances.
afi-all	(Optional) Specifies all address families.
ipv4	(Optional) Specifies IP Version 4 address prefixes.
ipv6	(Optional) Specifies IP Version 6 address prefixes.
unicast	(Optional) Specifies unicast address prefixes.
multicast	(Optional) Specifies multicast address prefixes.
safi-all	(Optional) Specifies unicast and multicast address prefixes.
detail	(Optional) Displays a detailed summary of the contents of the RIB, including the number of paths and some protocol-specific route attributes.
standby	(Optional) Displays standby information.

Command Default

If a **vrf** *vrf-name* is not specified, the contents of the RIB are displayed for the default IPv4 unicast VRF.

Command Modes

EXEC

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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **show route summary** command to display information about routes in the routing information base.

When a route summary is needed frequently—for instance, in a polling situation—use the **show route summary** command without the **detail** keyword. The **detail** keyword is used less frequently for verification purposes, because it is much more expensive (in bandwidth), requiring a scan of the entire routing database.

The **topology** keyword must be accompanied by the **ipv4 multicast** keywords, except when the **afi-all** keyword or the **safi-all** keyword is specified.

Task ID

Task ID	Operations
rib	read

Examples

The following is sample output from the **show route summary** command:

```
RP/0/RSP0/CPU0:router# show route summary

Route Source    Routes      Backup      Deleted      Memory (bytes)
static          1           0           0            136
connected       2           1           0            408
local           3           0           0            408
ospf            1673        2           0            272
isis            2           0           0            272
Total           10          1           0            1496
```

This table explains fields in the output of the **show route summary** command.

Table 164: show route summary Field Descriptions

Field	Description
Route Source	Routing protocol name.
Routes	Number of selected routes that are present in the routing table for each route source.
Backup	Number of routes that are not selected (are backup to a selected route).
Deleted	Number of routes that have been marked for deletion in the RIB, but have not yet been purged.
Memory	Number of bytes allocated to maintain all routes for the particular route source.

The following is sample output from the **show route summary** command with the **detail** keyword:

```
RP/0/RSP0/CPU0:router# show route summary detail

Route Source    Active Route     Active Path     Backup Route     Backup Path
static          1                1               0                0
connected       2                2               1                1
local           3                3               0                0
isis            1                1               1                1
Level 1:        0                0               1                1
Level 2:        1                1               0                0
ospf 1673       6                12              0                0
Intra-Area:     3                6               0                0
Inter-Area:     3                6               0                0
External-1:     0                0               0                0
External-2:     0                0               0                0
bgp 100         10               20              4                8
External:       5                10              4                8
Internal:       5                10              0                0
local:          0                0               0                0
Total           7                7               2                2
```

This table explains fields in the output of the **show route summary detail** command.

Table 165: show route summary detail Field Descriptions

Field	Description
Route Source	Source of the route. Routing protocol name and type.
Active Route	Number of active routes present in the routing table for each route source.
Active Path	Number of active paths present in the routing table for each route source.
Backup Route	Number of routes that are backup to a selected route for each route source.
Backup Path	Number of paths that are backup to a selected path for each route source.

Related Commands

Command	Description
show route, on page 1500	Displays the current contents of the routing table.



RIP Commands

This module describes the commands used to configure and monitor the Routing Information Protocol (RIP).

For detailed information about RIP concepts, configuration tasks, and examples, see the *Implementing RIP on* Cisco ASR 9000 Series Router module in the *Routing Configuration Guide for Cisco ASR 9000 Series Routers*.

- authentication keychain mode (RIP), on page 1531
- auto-summary (RIP), on page 1533
- broadcast-for-v2, on page 1534
- clear rip , on page 1535
- clear rip database, on page 1537
- clear rip interface, on page 1539
- clear rip out-of-memory, on page 1541
- clear rip statistics, on page 1543
- default-information originate (RIP), on page 1545
- default-metric (RIP), on page 1546
- distance (RIP), on page 1548
- interface (RIP), on page 1551
- maximum-paths (RIP), on page 1553
- metric-zero-accept, on page 1554
- neighbor (RIP), on page 1555
- nsf (RIP), on page 1557
- output-delay, on page 1558
- passive-interface (RIP), on page 1559
- poison-reverse, on page 1561
- receive version, on page 1563
- redistribute (RIP), on page 1564
- router rip, on page 1568
- route-policy (RIP), on page 1570
- send version, on page 1572
- show protocols (RIP), on page 1573
- show rip, on page 1575
- show rip database, on page 1577
- show rip interface, on page 1580
- show rip statistics, on page 1583

- site-of-origin (RIP), on page 1585
- split-horizon disable (RIP), on page 1587
- timers basic, on page 1589
- validate-update-source disable, on page 1591
- vrf (RIP), on page 1592

authentication keychain mode (RIP)

To enable an authentication keychain mechanism on RIP interfaces, use the **authentication keychain mode** command in interface configuration mode or VRF-interface configuration mode. To disable authentication keychain configuration on RIP interfaces, use the **no** form of this command.

authentication keychain keychain_name mode { **md5** | **text** }
no authentication keychain keychain_name mode { **md5** | **text** }

Syntax Description

keychain-name	Specifies the name of the keychain configured using the keychain command. **Note** All keychains need to be configured in Cisco IOS XR keychain database using the keychain configuration commands described in *Implementing Keychain Management* module of *System Security Configuration Guide for Cisco ASR 9000 Series Routers*
md5	Specifies that the authentication keychain mode is keyed message digest (md5).
text	Specifies that the authentication keychain mode is clear text.

Command Default

Keychain authentication is disabled.

Command Modes

Interface configuration

VRF-interface configuration

Command History

Release	Modification
Release 4.0.0	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.

All keychains need to be configured in Cisco IOS XR keychain database using the keychain configuration commands described in *Implementing Keychain Management* module of *System Security Configuration Guide for Cisco ASR 9000 Series Routers*

Task ID

Task ID	Operation
rip	read, write

This example shows how to configure an authentication keychain in md5 mode on a RIP VRF interface:

```
Router#configure
```

```
Router(config)#router rip
Router(config-rip)#vrf vrf_rip_auth
Router(config-rip-vrf)#interface POS 0/6/0/0
Router(config-rip-vrf-if)#authentication keychain key1 mode md5
```

This example shows how to configure an authentication keychain in clear text mode on a RIP interface:

```
Router#configure
Router(config)#router rip
Router(config-rip)#interface POS 0/6/0/1
Router(config-rip-if)#authentication keychain key2 mode text
```

Related Commands

Command	Description
router rip, on page 1568	Configures a routing process and enters router configuration mode for a Routing Information Protocol (RIP) process.
vrf (RIP), on page 1592	Defines a VPN routing and forwarding (VRF) instance and enters VRF configuration mode. Refer *System Security Command Reference for Cisco ASR 9000 Series Routers* for complete command reference information.
key chain (key chain)	Creates or modifies a keychain. Refer *System Security Command Reference for Cisco ASR 9000 Series Routers* for complete command reference information.
key (key chain)	Creates or modifies a keychain key. Refer *System Security Command Reference for Cisco ASR 9000 Series Routers* for complete command reference information.
key-string (keychain)	Specifies text string for the key. Refer *System Security Command Reference for Cisco ASR 9000 Series Routers* for complete command reference information.

auto-summary (RIP)

To enable the automatic summarization of subnet routes into network-level routes, use the **auto-summary** command in the appropriate configuration mode. To disable this function and send subprefix routing information across classful network boundaries, use the **no** form of this command.

auto-summary
no auto-summary

Syntax Description

This command has no arguments or keywords.

Command Default

Disabled

Command Modes

Router configuration

VRF 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **auto-summary** command to turn on route summarization. Route summarization reduces the amount of routing information in the routing tables.

Disable automatic summarization if you must perform routing between disconnected subnets. When automatic summarization is off, subnets are advertised. Automatic summarization is disabled by default.

Task ID

Task ID	Operations
rip	read, write

Examples

The following example shows how to turn on RIP auto-summarization:

```
RP/0/RSP0/CPU0:router(config)# router rip
RP/0/RSP0/CPU0:router(config-rip)# auto-summary
```

Related Commands

Command	Description
vrf (RIP), on page 1592	Defines a VRF instance and enters VRF configuration mode.

broadcast-for-v2

To send Routing Information Protocol (RIP) Version 2 output packets to a broadcast address, use the **broadcast-for-v2** command in the appropriate configuration mode. To disable this feature, use the **no** form of this command.

broadcast-for-v2
no broadcast-for-v2

Syntax Description

This command has no arguments or keywords.

Command Default

RIPv2 output packets are not broadcasted.

Command Modes

Router configuration

VRF configuration

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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **broadcast-for-v2** command to broadcast RIP Version 2 broadcast updates to hosts that do not listen to multicasts. Version 2 updates (requests and responses) will be sent to the IP broadcast address 255.255.255.255 instead of the IP multicast address 244.0.0.9.

Task ID

Task ID	Operations
rip	read, write

Examples

The following example shows how to send RIP v2 output messages to a broadcast address for all RIP interfaces:

```
RP/0/RSP0/CPU0:router(config)# router rip
RP/0/RSP0/CPU0:router(config-rip)# broadcast-for-v2
```

clear rip

To clear VRF and interface-related information for Routing Information Protocol (RIP) such as database entries and statistics, use the **clear rip** command in EXEC configuration mode.

clear rip [**vrf** { *name* | **all** }] [**database** | **out-of-memory** | **statistics**] [**interface** *type number*]

Syntax Description

vrf {*name*|**all**}	(Optional) Clears RIP-associated VRF information. You can specify a VRF, or clear all VRF information.
database	(Optional) Clears database entries from the RIP topology table.
out-of-memory	(Optional) Clears RIP out-of-memory state information.
statistics	(Optional) Clears RIP statistics information.
interface *type number*	(Optional) Clears RIP interface information.

Command Default

No default behavior or values

Command Modes

EXEC configuration

Command History

Release	Modification
Release 7.4.1	This command was introduced.

Usage Guidelines

RIP VRFs and interfaces that are forcibly deactivated by the software because of a severe memory state are not activated again until the out-of-memory state is cleared by using the **clear rip** , **clear rip interface** , or **clear rip out-of-memory** command. Use the **clear rip out-of-memory** command to clear the out-of-memory state completely and allow the RIP process to force the VRF or interface to shut down.

When the router begins to run out of memory, the RIP process can transition through different memory states defined as Normal, Minor, Severe, and Critical:

- In Normal state, RIP VRFs and interfaces function normally.
- In Minor state, RIP VRFs and interfaces that are currently active are allowed to remain active. VRFs and interfaces that are not currently active are not allowed to become active until the RIP process transitions to Normal state.
- In Severe state, a few VRFs and interfaces are forcibly brought down periodically until the RIP process transitions to another state.
- In Critical state, the RIP process is forcibly shut down.

Examples

The following example shows how to clear all database, interface, and VRF entries in RIP:

```
Router# clear rip vrf all
```

The following example shows how to clear only database entries from the topology table for the GigabitEthernet 0/1/0/0 interface:

```
Router# clear rip database interface GigabitEthernet 0/1/0/0
```

clear rip database

To clear only database entries from the Routing Information Protocol (RIP) topology table, use the **clear rip database** command in EXEC configuration mode.

clear rip [**vrf** {*vrf* | **all**}] **database** [**interface** *type interface-path-id*]

Syntax Description

vrf { *vrf* | **all** }	(Optional) Specifies a particular VPN routing and forwarding (VRF) instance or all VRF instances.
interface	(Optional) Specifies the interface to clear topology entries.
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

No default behavior or values

Command Modes

EXEC 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 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
rip	read, write

Examples

The following example shows how to clear only database entries from the topology table for the GigabitEthernet 0/1/0/0 interface:

```
RP/0/RSP0/CPU0:router# clear rip database interface GigabitEthernet 0/1/0/0
```

Related Commands

Command	Description
show rip statistics, on page 1583	Displays database and interface entry information for the RIP process.

clear rip interface

To clear interface-related information for Routing Information Protocol (RIP) such as database entries and statistics, use the **clear rip interface** command in EXEC configuration mode.

clear rip [**vrf** {**vrf** | **all**}] **interface** *type interface-path-id*

Syntax Description

vrf { *vrf* | **all** }	(Optional) Specifies a particular VPN routing and forwarding (VRF) instance or all VRF instances.
interface	Specifies the interface to clear topology entries.
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

No default behavior or values

Command Modes

EXEC 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

RIP VRFs and interfaces that are forcibly deactivated by the software because of a severe memory state are not activated again until after out-of-memory state is cleared by using the **clear rip** , **clear rip interface** or **c lear rip out-of-memory** command.

Task ID

Task ID	Operations
rip	read, write

Examples

The following example shows how to clear all interface-related data such as routes and statistics from the GigabitEthernet 0/1/0/0 interface:

```
RP/0/RSP0/CPU0:router# clear rip interface GigabitEthernet 0/1/0/0
```

Related Commands

Command	Description
#unique_758	Clears VRF and interface-related information for RIP such as database entries and statistics.
clear rip out-of-memory, on page 1541	Clears the out-of-memory state for RIP.

clear rip out-of-memory

To clear the out-of-memory state for Routing Information Protocol (RIP), use the **clear rip out-of-memory** command in EXEC configuration mode.

clear rip [**vrf** {*vrf* | **all**}] **out-of-memory** [**interface** *type interface-path-id*]

Syntax Description

vrf { *vrf* | **all** }	(Optional) Specifies a particular VPN routing and forwarding (VRF) instance or all VRF instances.
interface	(Optional) Specifies the interface to clear topology entries.
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

No default behavior or values

Command Modes

EXEC 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **clear rip out-of-memory** command, to clear the out-of-memory state completely and allow the RIP process to force the VRF or interface to shut down.

When the router begins to run out of memory, the RIP process can transition through different memory states defined as Normal, Minor, Severe, and Critical:

- In Normal state, RIP VRFs and interfaces function normally.
- In Minor state, RIP VRFs and interfaces that are currently active are allowed to remain active. VRFs and interfaces that are not currently active are not allowed to become active until the RIP process transitions to Normal state.
- In Severe state, a few VRFs and interfaces are forcibly brought down periodically until the RIP process transitions to another state.
- In Critical state, the RIP process is forcibly shut down.

VRFs and interfaces that are forcibly shut down in Severe state are not automatically activated when the RIP process transitions to Minor or Normal state. When a VRF or interface is forcibly brought down while in

Severe state, the **clear rip** , **clear rip interface** or **clear rip out-of-memory** command clears the Forced Down state and reactivates the VRF or interface.

The **show rip** and **show rip interface** commands allow you to view the current out-of-memory state.

Task ID

Task ID	Operations
rip	read, write

Examples

The following example shows how to clear the out-of-memory state for a RIP process:

```
RP/0/RSP0/CPU0:router# clear rip out-of-memory
```

Related Commands

#unique_758	Clears VRF and interface-related information for RIP such as database entries and statistics.
clear rip interface, on page 1539	Clears interface-related information for RIP such as database entries and statistics.
show rip, on page 1575	Displays configuration and status of RIP.
show rip interface, on page 1580	Displays interface entry information from the RIP topology table.

clear rip statistics

To clear the Routing Information Protocol (RIP) statistics, use the **clear rip statistics** command in EXEC configuration mode.

clear rip [**vrf** {**vrf** | **all**}] **statistics** [**interface** *type interface-path-id*]

Syntax Description

vrf { *vrf* | **all** }	(Optional) Specifies a particular VPN routing and forwarding (VRF) instance or all VRF instances.
interface	(Optional) Specifies the interface from which to clear topology entries.
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

No default behavior or values

Command Modes

EXEC 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 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
rip	read, write

Examples

The following example shows how to clear all RIP statistics:

```
RP/0/RSP0/CPU0:router# clear rip statistics
```

Related Commands

Command	Description
show rip statistics, on page 1583	Displays database and interface entry information for the RIP process.

default-information originate (RIP)

To generate a default route into Routing Information Protocol (RIP), use the **default-information originate** command in the appropriate configuration mode. To disable a default route into RIP, use the **no** form of this command.

default-information originate [**route-policy** *name*]
no default-information originate

Syntax Description

route-policy	*name*	Route policy name that indicates criteria for the default route.

Command Default

This command is disabled by default.

Command Modes

Router configuration

VRF configuration

Command History

Release	Modification
Release 3.7.2	This command was introduced.
Release 7.5.2	Support for IPv6 capabilities 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
rip	read, write

Examples

The following example shows how to originate a default route in RIP updates based on the result of running the route policy on the routing table:

```
Router(config)# router rip
Router(config-rip)# default-information originate route-policy policy1
```

Related Commands

Command	Description
route-policy (RIP), on page 1570	Applies a routing policy to updates advertised to or received from a RIP neighbor.

default-metric (RIP)

To set default metric values for routes redistributed from other protocols into Routing Information Protocol (RIP), use the **default-metric** command in the appropriate configuration mode. To return to the default state, use the **no** form of this command.

default-metric *number-value*
no default-metric

Syntax Description

number-value	Default metric value. Range is 1 to 15.

Command Default

Default metrics are not set.

Command Modes

Router configuration

VRF 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **default-metric** command with the **redistribute** command to cause RIP to use the same metric value for all redistributed routes. A default metric helps solve the problem of redistributing routes with incompatible metrics by providing a reasonable substitute and enables redistribution to proceed. If you want to set different metrics for other redistributed protocols, use the **route-policy** option in the **redistribute** command.

The RIP metric used for redistributed routes is determined by the route policy. If a route policy is not configured or the route policy does not set the RIP metric, the metric is determined based on the redistributed protocol. For VPNv4 routes redistributed by BGP, the RIP metric set at the remote PE router is used, if valid.

In all other cases (BGP, IS-IS, OSPF, EIGRP, connected, static), the metric set by the **default-metric** command is used. If a valid metric cannot be determined, then redistribution does not happen.

Task ID

Task ID	Operations
rip	read, write

Examples

The following example shows how a router in autonomous system 109 uses both the RIP and the Open Shortest Path First (OSPF) routing protocols. The example advertises OSPF-derived routes using RIP and assigns the OSPF-derived routes a RIP metric of 10:

```
RP/0/RSP0/CPU0:router(config)# router rip
```

```
RP/0/RSP0/CPU0:router(config-rip)# vrf vpn-1
RP/0/RSP0/CPU0:router(config-rip-vrf)# default-metric 10
RP/0/RSP0/CPU0:router(config-rip-vrf)# redistribute ospf 109
```

Related Commands

Command	Description
redistribute (RIP), on page 1564	Redistributes routes from one routing domain into RIP.

distance (RIP)

To define the administrative distance assigned to routes discovered by the Routing Information Protocol (RIP), use the **distance admin-distance** command in the appropriate configuration mode. To remove the distance definition from the configuration file and restore the system to its default condition, use the **no** form of this command.

distance *admin-distance* [*prefix prefix-length* | *prefix mask*]
no distance *admin-distance*

Syntax Description

admin-distance	Administrative distance to be assigned to RIP routes. Range is 0 to 255.
prefix	(Optional) Network IP address about which routing information should be displayed.
prefix-length	(Optional) The *prefix-length* argument specifies the length of the IPv4 or IPv6 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. Range is 0 to 32 for IPv4 addresses.
mask	(Optional) Network mask specified in either of two ways: • Network mask can be a four-part, dotted decimal address. For example, 255.0.0.0 indicates that each bit equal to 1 means the corresponding address bit is a network address. • Network mask can be indicated as a slash (/) and number. For example, /8 shows that the first 8 bits of the mask are ones, and the corresponding bits of the address are the network address.

Command Default

admin-distance : 120

Command Modes

Router configuration

VRF configuration

Command History

Release	Modification
Release 3.7.2	This command was introduced.
Release 7.5.2	Support for IPv6 capabilities 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.

Use the **distance** command to change the preference of RIP routes over other protocol routes. When administrative distance and redistribution features are used to together, routing behavior may be affected for routes accepted from and advertised to RIP neighbors.

Numerically, an administrative distance is an integer from 0 to 255. In general, the higher the value, the lower the trust rating. An administrative distance of 255 means that the routing information source cannot be trusted at all and should be ignored.

The order in which you enter distance commands can affect the assigned administrative distances in unexpected ways.

This table lists default administrative distances.

Table 166: Default Administrative Distances of Routing Protocols

Routing Protocol	Administrative Distance Value
Connected interface	0
Static route out an interface	0
Static route to next-hop	1
EIGRP Summary Route	5
External BGP	20
Internal EIGRP	90
OSPF	110
IS-IS	115
RIP Versions 1 and 2	120
External EIGRP	170
Internal BGP	200
Unknown	255

Task ID

Task ID	Operations
rip	read, write

Examples

The following example shows how to set the administrative distance for a particular IPv4 prefix:

```
Router(config)# router rip
Router(config-rip)# distance 85 192.168.10.0/24
```

Examples

The following example shows how to set the administrative distance for a particular IPv4 prefix:

```
Router(config)# router rip
Router(config-rip)# distance 85 2001:DB8::1/32
```

Related Commands

Command	Description
redistribute (RIP), on page 1564	Redistributes routes from one routing domain into RIP.

interface (RIP)

To define the interfaces on which the Routing Information Protocol (RIP) runs and enter interface configuration mode, use the **interface** command in router configuration mode. To disable RIP routing for interfaces, use the **no** form of this command.

interface *type interface-path-id*
no interface *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 a 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

When you do not specify this command in configuration mode, RIP routing for interfaces is not enabled.

Command Modes

Router configuration

VRF 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **interface** command to associate a specific interface with a RIP process. The interface remains associated with the process even when the IPv4 address of the interface changes.

This command places the router in interface configuration mode, from which you can configure interface-specific settings. Commands configured under this mode (such as the broadcast-for-v2, on page 1534 command) are automatically bound to that interface.

Task ID

Task ID	Operations
rip	read, write

Examples

The following example shows how to enter interface configuration mode for a RIP process, and send RIP Version 2 messages to the broadcast address on the GigabitEthernet interface 0/1/0/0:

```
RP/0/RSP0/CPU0:router(config)# router rip
RP/0/RSP0/CPU0:router(config-rip)# interface GigabitEthernet 0/1/0/0
RP/0/RSP0/CPU0:router(config-rip-if)# ?

  broadcast-for-v2    Specify broadcast address for RIP v2 output packet
  commit              Commit the configuration changes to running
  describe            Describe a command without taking real actions
  do                  Run an exec command
  exit                Exit from this submode
  metric-zero-accept  Accept rip update with metric 0 to compensate a common bug
  no                  Negate a command or set its defaults
  passive-interface   Suppress routing updates on this interface
  poison-reverse      Enable poison reverse
  receive             Advertisement reception
  route-policy        Apply route policy to routing updates
  send                Advertisement transmission
  show                Show contents of configuration
  site-of-origin      SOO community for prefixes learned over this interface
  split-horizon       Disable split horizon
RP/0/RSP0/CPU0:router(config-rip-if)# broadcast-for-v2
```

Related Commands

Command	Description
broadcast-for-v2, on page 1534	Sends RIP Version 2 output packets to a broadcast address.

maximum-paths (RIP)

To configure the maximum number of equal cost parallel routes that the Routing Information Protocol (RIP) will install into the routing table, use the **maximum-paths** command in the appropriate configuration mode. To remove the **maximum-paths** command from the configuration file and restore the system to its default condition with respect to RIP, use the **no** form of this command.

maximum-paths *maximum*
no maximum-paths

Syntax Description

maximum	Maximum number of parallel routes that RIP can install in a routing table. Range is 1 to 32..

Command Default

4 paths

Command Modes

Router configuration

VRF configuration

Command History

Release	Modification
Release 3.7.2	This command was introduced.
Release 7.5.2	Support for IPv6 capabilities 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
rip	read, write

Examples

The following example shows how to allow a maximum of 16 equal cost paths to a destination:

```
RP/0/RSP0/CPU0:router(config)# router rip
RP/0/RSP0/CPU0:router(config-rip)# maximum-paths 16
```

metric-zero-accept

To allow RIP to accept routing entries from RIP updates with a metric set to zero (0), use the **metric-zero-accept** command in interface configuration mode. To remove the **metric-zero-accept** command from the configuration file and restore the system to its default condition with respect to RIP, use the **no** form of this command.

metric-zero-accept
no metric-zero-accept

Syntax Description

This command has no arguments or keywords.

Command Default

RIP routes received with a metric of zero (0) are ignored.

Command Modes

Interface configuration

Command History

Release	Modification
Release 3.7.2	This command was introduced.
Release 7.5.2	Support for IPv6 capabilities 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.

After the **metric-zero-accept** command is configured on routing entries from RIP updates, RIP accepts these routes and then sets the metric to one (1).

Task ID

Task ID	Operations
rip	read, write

Examples

The following example shows how to set the RIP interface to accept metric zero on routing entries:

```
Router(config)# router rip
Router(config-rip)# interface GigabitEthernet 0/1/0/0
Router(config-rip-if)# metro-zero-accept
```

neighbor (RIP)

To define a neighboring router with which to exchange Routing Information Protocol (RIP) information, use the **neighbor** command in the appropriate configuration mode. To remove an entry, use the **no** form of this command.

neighbor { *ipv4-address* | *ipv6-address*
no neighbor { *ipv4-address* | *ipv6-address*

Syntax Description

ipv4-address	IPv4 address of a peer router with which routing information is exchanged.
ipv6-address	IPv6 link-local address of a peer router with which routing information is exchanged.

Command Default

No neighboring routers are defined.

Command Modes

Router configuration

VRF configuration

Command History

Release	Modification
Release 3.7.2	This command was introduced.
Release 7.5.2	Support for IPv6 capabilities 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.

Use the **neighbor** command to permit the point-to-point (nonbroadcast) exchange of routing information. When the **neighbor** command is used in combination with the **passive-interface** command in router configuration mode, routing information can be exchanged between a subset of routers and access servers on a LAN.

Multiple **neighbor** commands can be used to specify additional neighbors or peers.

Task ID

Task ID	Operations
rip	read, write

Examples

The following example shows how to permit the sending of RIP updates to specific neighbors. One copy of the update is generated per neighbor:

```
RP/0/RSP0/CPU0:router(config)# router rip
RP/0/RSP0/CPU0:router(config-rip)# vrf vpn-1
RP/0/RSP0/CPU0:router(config-rip-vrf)# neighbor 172.16.1.2
```

Related Commands

Command	Description
passive-interface (RIP), on page 1559	Suppresses the sending of RIP updates on an interface.

nsf (RIP)

To configure nonstop forwarding (NSF) on Routing Information Protocol (RIP) routes after a RIP process shutdown or restart, use the **nsf** command in the appropriate configuration mode. To remove this command from the configuration file and restore the system to its default condition, use the **no** form of this command.

nsf
no nsf

Syntax Description

This command has no arguments or keywords.

Command Default

NSF is disabled.

Command Modes

Router configuration

VRF 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

When you use the **nsf** command, NSF lifetime is automatically set to two times the update time (with a minimum value of 60 seconds). The RIP process must reconverge within this time. If the convergence exceeds the NSF lifetime, routes are purged from the Routing Information Base (RIB) and NSF may fail.

Task ID

Task ID	Operations
rip	read, write

Examples

The following example shows how to configure RIP NSF:

```
RP/0/RSP0/CPU0:router(config)# router rip
RP/0/RSP0/CPU0:router(config-rip)# vrf vpn-1
RP/0/RSP0/CPU0:router(config-rip-vrf)# nsf
```

output-delay

To change the interpacket delay for Routing Information Protocol (RIP) updates sent, use the **output-delay** command in the appropriate configuration mode. To remove the delay, use the **no** form of this command.

output-delay *delay*
no output-delay *delay*

Syntax Description

delay	Delay (in milliseconds) between consecutive packets in a multiple-packet RIP update. The range is from 8 to 50.

Command Default

The default is no delay.

Command Modes

Router configuration

VRF 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **output-delay** command if you are sending at high speed to a low-speed router that might not be able to receive at the high speed. Configuring this command helps prevent the routing table from losing information.

Task ID

Task ID	Operations
rip	read, write

Examples

The following example shows how to set the interpacket delay to 10 milliseconds:

```
RP/0/RSP0/CPU0:router(config)# router rip
RP/0/RSP0/CPU0:router(config-rip)# vrf vpn-1
RP/0/RSP0/CPU0:router(config-rip-vrf)# output-delay 10
```

passive-interface (RIP)

To suppress the sending of Routing Information Protocol (RIP) updates on an interface, use the **passive-interface** command in interface configuration mode. To unsuppress updates, use the **no** form of this command.

passive-interface
no passive-interface

Syntax Description

This command has no arguments or keywords.

Command Default

RIP updates are sent on the interface.

Command Modes

Interface configuration

Command History

Release	Modification
Release 3.7.2	This command was introduced.
Release 7.5.2	Support for IPv6 capabilities 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.

While RIP stops sending routing updates to the multicast (or broadcast) address on a passive interface, RIP continues to receive and process routing updates from its neighbors on that interface.

Task ID

Task ID	Operations
rip	read, write

Examples

The following example shows that GigabitEthernet interface 0/1/0/0 stops multicasting (or broadcasting) RIP updates while continuing to receive RIP updates normally. GigabitEthernet interface 0/1/0/3 sends and receives updates normally. Also RIP updates are unicast to neighbor 172.168.1.2 over the appropriate interface:

```
Router(config)# router rip
Router(config-rip)# neighbor 172.16.1.2
Router(config-rip)# interface GigabitEthernet 0/1/0/0
Router(config-rip-if)# passive-interface
Router(config-rip-if)# exit
Router(config-rip)# interface GigabitEthernet 0/1/0/3
Router(config-rip-if)# exit
```

Examples

The following example shows that GigabitEthernet interface 0/1/0/0 stops multicasting (or broadcasting) RIP updates while continuing to receive RIP updates normally. GigabitEthernet interface 0/1/0/3 sends and receives updates normally. Also RIP updates are unicast to neighbor 2001:DB8::1 over the appropriate interface:

```
Router(config)# router rip
Router(config-rip)# neighbor 2001:DB8::1
Router(config-rip)# interface GigabitEthernet 0/1/0/0
Router(config-rip-if)# passive-interface
Router(config-rip-if)# exit
Router(config-rip)# interface GigabitEthernet 0/1/0/3
Router(config-rip-if)# exit
```

Related Commands

Command	Description
neighbor (RIP), on page 1555	Defines a neighboring router with which to exchange RIP protocol information.

poison-reverse

To enable poison reverse processing of Routing Information Protocol (RIP) router updates, use the **poison-reverse** command in interface configuration mode. To disable poison reverse processing of RIP updates, use the **no** form of this command.

poison-reverse
no poison-reverse

Syntax Description

This command has no arguments or keywords.

Command Default

Poison reverse processing is disabled.

Command Modes

Interface configuration

Command History

Release	Modification
Release 3.7.2	This command was introduced.
Release 7.5.2	Support for IPv6 capabilities 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.

Route poisoning prevents routing loops by communicating to other routers that a route is no longer reachable, effectively removing these routes from other router's routing tables. The system default, **split horizon** , provides that routes learned through RIP are not advertised from the interface over which they were learned.

The **poison-reverse** command enables poison reverse processing of RIP router updates. A router that receives route poisoning information sends the poisoning information back to the sending router, a process called poison reverse. This process ensures that all routers on the same interface have received the poisoned route information.

If both **poison-** reverse and split horizon are configured, then simple split horizon behavior (suppression of routes from the interface over which they were learned) is replaced by poison reverse behavior. If split horizon is disabled, the poison reverse configuration is ignored.

Task ID

Task ID	Operations
rip	read, write

Examples

The following example shows how to enable poison reverse processing for an interface running RIP:

```
Router(config)# router rip
Router(config-rip)# interface GigabitEthernet 0/1/0/0
Router(config-rip-if)# poison-reverse
```

Related Commands

Command	Description
split-horizon disable (RIP), on page 1587	Disables the split horizon mechanism.

receive version

To configure the Routing Information Protocol (RIP) interface to accept version-specific packets, use the **receive version** command in interface configuration mode. To revert to the default setting, use the **no** form of this command.

receive version {**1** | **2** | **1 2**}
no receive version {**1** | **2** | **1 2**}

Syntax Description

1	Version 1 packets.
2	Version 2 packets.
1 2	Both versions 1 and 2 packets.

Command Default

Version 2

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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **receive version** command to override the default behavior of RIP. This command applies only to the interface being configured.

Task ID

Task ID	Operations
rip	read, write

Examples

The following example shows how to configure an interface to accept both RIP Version 1 and 2 packets:

```
RP/0/RSP0/CPU0:router(config)# router rip
RP/0/RSP0/CPU0:router(config-rip)# interface GigabitEthernet 0/1/0/0
RP/0/RSP0/CPU0:router(config-rip-if)# receive version 1 2
```

Related Commands

Command	Description
send version, on page 1572	Configures the RIP interface to send version specific packets.

redistribute (RIP)

To redistribute routes from another routing domain into Routing Information Protocol (RIP), use the **redistribute** command in the appropriate configuration mode. To remove the **redistribute** command from the configuration file and restore the system to its default condition in which the software does not redistribute routes, use the **no** form of this command.

Border Gateway Protocol (BGP)
redistribute bgp *process-id* [**route-policy** *name*] [**external** | **internal** | **local**]
no redistribute bgp *process-id*

Connected Interface Routes
redistribute connected [**route-policy** *name*]
no redistribute connected

Enhanced Interior Gateway Routing Protocol (EIGRP)
redistribute eigrp *process-id* [**route-policy** *name*]
no redistribute eigrp *process-id*

Intermediate System-to-Intermediate System (ISIS)
redistribute isis *process-id* [**route-policy** *name*] [**level-1** | **level-1-2** | **level-2**]
no redistribute isis *process-id*

Open Shortest Path First (OSPF)
redistribute ospf *process-id* [**route-policy** *name*] [**match** {**external** [**1** | **2**] | **internal** | **nssa-external** [**1** | **2**]}]
no redistribute ospf *process-id*

IP Static Routes
redistribute static [**route-policy** *name*]
no redistribute static

Syntax Description

bgp	Distributes routes from the BGP protocol.

process-id	• For the **bgp** keyword:
	  • 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.
	• For the **eigrp** keyword, an EIGRP instance name from which routes are to be redistributed. The value takes the form of a string. A decimal number can be entered, but it is stored internally as a string.
	• For the **isis** keyword, an IS-IS instance name from which routes are to be redistributed. The value takes the form of a string. A decimal number can be entered, but it is stored internally as a string.
	• For the **ospf** keyword, an OSPF instance name from which routes are to be redistributed. The value takes the form of a string. A decimal number can be entered, but it is stored internally as a string.

external	(Optional) Specifies BGP external routes only.
internal	(Optional) Specifies BGP internal routes only.
local	(Optional) Specifies BGP local routes only.
route-policy *name*	(Optional) Specifies the identifier of a configured policy. A policy is used to filter the importation of routes from this source routing protocol to RIP.
level-1	(Optional) Redistributes Level 1 IS-IS routes into other routing protocols independently.
level-1-2	(Optional) Distributes both Level 1 and Level 2 IS-IS routes into other routing protocols.
level-2	(Optional) Distributes Level 2 IS-IS routes into other routing protocols independently.
[**match** { **external** [**1** | **2**] | **internal** | **nssa-external** [**1** | **2**]]} [**route-policy** *name*]	(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. If no match is specified, the default is no filtering.
static	Redistributes IP static routes.

Command Default

Route redistribution is disabled.

Command Modes

Router configuration

VRF configuration

Command History

Release	Modification
Release 3.7.2	This command was introduced.
Release 3.9.0	Asplain format for 4-byte Autonomous system numbers notation was supported.
Release 7.5.2	Support for IPv6 capabilities 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.

Note When you are redistributing routes (into RIP) 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.

Redistributed routing information may be filtered by the **route-policy** *name* keyword and argument. This filtering ensures that only those routes intended by the administrator are redistributed by RIP.

The RIP metric used for redistributed routes is determined by the route policy. If a route policy is not configured or the route policy does not set the RIP metric, the metric is determined based on the redistributed protocol. For VPNv4 routes redistributed by BGP, the RIP metric set at the remote PE router is used, if valid.

In all other cases (BGP, IS-IS, OSPF, EIGRP, connected, static), the metric set by the **default-metric** command is used. If a valid metric cannot be determined, then redistribution does not happen.

For information about routing policies, see the *Routing Policy Commands on* Cisco ASR 9000 Series Router module of the *Cisco ASR 9000 Series Aggregation Services Router Routing Command Reference*.

Task ID

Task ID	Operations
rip	read, write

Examples

The following example shows how to cause BGP routes to be redistributed into a RIP process:

```
RP/0/RSP0/CPU0:router(config)# router rip
RP/0/RSP0/CPU0:router(config-rip)# redistribute bgp 100
```

Related Commands

Command	Description
default-metric (RIP), on page 1546	Sets default metric values for routes redistributed from other protocols into RIP.

router rip

To configure a routing process and enter router configuration mode for a Routing Information Protocol (RIP) process, use the **router rip** command in global configuration mode. To turn off the RIP routing process, use the **no** form of this command.

router rip
no router rip

Syntax Description

This command has no arguments or keywords.

Command Default

No router process is defined.

Command Modes

Global 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 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
rip	read, write

Examples

The following example shows how to configure a router process for RIP:

```
RP/0/RSP0/CPU0:router(config)# router rip
RP/0/RSP0/CPU0:router(config-rip)#
```

The following example shows how to enter router configuration mode for RIP and identify commands that can be issued from that mode.

```
RP/0/RSP0/CPU0:router(config)# router rip
RP/0/RSP0/CPU0:router(config-rip)# ?

  auto-summary            Enable automatic network number summarization
  broadcast-for-v2        Send RIP v2 output packets to broadcast address
  commit                  Commit the configuration changes to running
  default-information     Control distribution of default information
  default-metric          Set metric of redistributed routes
  describe                Describe a command without taking real actions
  distance                Define an administrative distance
  do                      Run an exec command
```

```
  exit                    Exit from this submode
  interface               Enter the RIP interface configuration submode
  maximum-paths           Maximum number of paths allowed per route
  neighbor                Specify a neighbor router
  no                      Negate a command or set its defaults
  nsf                     Enable Cisco Non Stop Forwarding
  output-delay            Interpacket delay for RIP updates
  redistribute            Redistribute information from another routing protocol
  route-policy            Apply route policy to routing updates
  show                    Show contents of configuration
  timers                  Adjust routing timers
  validate-update-source  Validate source address of routing updates
  vrf                     Enter the RIP vrf configuration submode
RP/0/RSP0/CPU0:router(config-rip)#
```

route-policy (RIP)

To apply a routing policy to updates advertised to or received from a Routing Information Protocol (RIP) neighbor, use the **route-policy** command in the appropriate configuration mode. To disable applying routing policy to updates, use the **no** form of this command.

route-policy *name* {**in** | **out**}
no route-policy *name* {**in** | **out**}

Syntax Description

name	Name of route policy.
in	Applies policy to inbound routes.
out	Applies policy to outbound routes.

Command Default

No policy is applied.

Command Modes

Router configuration

VRF configuration

Interface configuration

Command History

Release	Modification
Release 3.7.2	This command was introduced.
Release 7.5.2	Support for IPv6 capabilities 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.

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.

Note If a route policy is configured both on the interface and on the VRF, the interface route policy is applied.

Task ID

Task ID	Operations
rip	read, write

Examples

The following example shows how to filter routing updates received on an interface:

```
Router(config)# router rip
Router(config-rip)# interface GigabitEthernet 0/1/0/0
Router(config-rip-if)# route-policy updpol-1 in
```

send version

To configure the Routing Information Protocol (RIP) interface to send version specific packets, use the **send version** command in interface configuration mode. To revert to the default setting, use the **no** form of this command.

send version {**1** | **2** | **1 2**}
no send version {**1** | **2** | **1 2**}

Syntax Description

1	Version 1 packets.
2	Version 2 packets.
1 2	Both Version 1 and Version 2 packets.

Command Default

Version 2

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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **send version** command to override the default behavior of RIP. This command applies only to the interface being configured.

Task ID

Task ID	Operations
rip	read, write

Examples

The following example shows how to configure an interface to send only RIP Version 2 packets:

```
RP/0/RSP0/CPU0:router(config)# router rip
RP/0/RSP0/CPU0:router(config-rip)# interface GigabitEthernet 0/1/0/0
RP/0/RSP0/CPU0:router(config-rip-if)# send version 2
```

Related Commands

Command	Description
receive version, on page 1563	Configures the RIP interface to accept version-specific packets.

show protocols (RIP)

To display information about the Routing Information Protocol (RIP) process configuration, use the **show protocols** command in EXEC mode.

show protocols [**ipv4** | **afi-all**] [**all***protocol*] [**default-context** | [**vrf** {*vrf-name* | **all**}]] [**private**]

Syntax Description

ipv4	(Optional) Specifies an IPv4 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 **eigrp** , **bgp** , **isis** , **ospf** , and **rip** .
default-context	(Optional) Displays default context information. This keyword is available when the **eigrp** or **rip** protocol is specified.
vrf *vrf-name*	(Optional) Displays VPN routing and forwarding (VRF) information for the specified process. This keyword is available when the **eigrp** or **rip** protocol is specified.
private	(Optional) Displays private EIGRP data. This keyword is available when the **eigrp** protocol is specified.

Command Modes

EXEC

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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

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 summarizes the important characteristics of the running protocol, and command output varies depending on the specific protocol selected.

For RIP, the command output lists the instance number, default AS context, router ID, default networks, distance, maximum paths, and so on.

Task ID

Task ID	Operations
RIP	read

Examples

The following is sample output from the **show protocols rip** command:

```
RP/0/RSP0/CPU0:router# show protocols rip
Routing Protocol RIP
  2 VRFs (including default) configured, 2 active
  25 routes, 16 paths have been allocated
  Current OOM state is "Normal"
  UDP socket descriptor is 37
    VRF              Active   If-config If-active Routes     Paths      Updates
    default          Active   3         3         11         7              30s
```

This table describes the significant fields shown in the display.

Table 167: show protocols Field Descriptions

Field	Description
VRFs configured	Number of VRFs configured.
VRFs active	Number of active VRFs.
Routes	Number of allocated routes.
Paths	Number of allocated paths.
OOM state	Current out-of-memory state of RIP process.
UDP socket	Current UDP socket descriptor value.

show rip

To display configuration and status of Routing Information Protocol (RIP), use the **show rip** command in EXEC mode.

show rip [**vrf** {*vrf-name* | **all**}]

Syntax Description

vrf { *vrf* | **all** }	(Optional) Specifies a particular VPN routing and forwarding (VRF) instance or all VRF instances.

Command Default

No default behavior or values

Command Modes

EXEC

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 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
rip	read

Examples

The following example shows sample output from the **show rip** command:

```
RP/0/RSP0/CPU0:router# show rip

RIP config:
Active?:                    Yes
Added to socket?:           Yes
Out-of-memory state:        Normal
Version:                    2
Default metric:             Not set
Maximum paths:              4
Auto summarize?:            No
Broadcast for V2?:          No
Packet source validation?:  Yes
NSF:                        Disabled
Timers: Update:             30 seconds (25 seconds until next update)
        Invalid:            180 seconds
        Holddown:           180 seconds
        Flush:              240 seconds
```

This table describes the significant fields shown in the display.

Table 168: show rip Field Descriptions

Field	Description
Active?	Active state setting.
Added to socket?	Multicast group setting on RIP configured interfaces. If yes, updates are received on these interfaces.
Out-of-memory state	Out-of-memory state for RIP can be one of the following: Normal, Minor, Severe, or Critical.
Version	Version number is 2.
Default metric	Default metric value, if configured. Otherwise Not set.
Maximum paths	Number of maximum paths allowed per RIP route.
Auto summarize?	Auto-summarize state setting.
Broadcast for V2?	RIP Version 2 broadcast setting.
Packet source validation?	Validation setting for the source IP address of incoming routing updates to RIP.
Timers	RIP network timer settings.

show rip database

To display database entry information from the Routing Information Protocol (RIP) topology table, use the **show rip database** command in EXEC mode.

show rip [**vrf** { **IPv4** | **IPv6** } { *vrf-name* | **all** }] **database** [*prefix prefix-length* | *prefix mask*]

Syntax Description

IPv4 | **IPv6**	Specifies IPv4 or IPv6.
vrf { *vrf* | **all** }	(Optional) Specifies a particular VPN routing and forwarding (VRF) instance or all VRF instances.
prefix	(Optional) Network IP address about which routing information should be displayed.
prefix-length	(Optional) The *prefix-length* argument specifies the length of the IP 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 between must precede the decimal value. Range is 0 to 32 for IPv4 addresses.
prefix-mask	(Optional) Network mask specified in either of two ways: • Network mask can be a four-part, dotted decimal address. For example, 255.0.0.0 indicates that each bit equal to 1 means the corresponding address bit is a network address. • Network mask can be indicated as a slash (/) and number. For example, /8 shows that the first 8 bits of the mask are ones, and the corresponding bits of the address are the network address.

Command Default

No default behavior or values

Command Modes

EXEC

Command History

Release	Modification
Release 3.7.2	This command was introduced.
Release 7.5.2	Support for IPv6 capabilities 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.

Summary address entries appear in the database only if relevant child routes are summarized. When the last child route for a summary address becomes invalid, the summary address is also removed from the routing table.

Task ID

Task ID	Operations
rip	read

Examples

The following is sample output from the **show rip ipv4 database** command:

```
Router# show rip ipv4 database

IPv4 Routes held in RIP's topology database:
    10.0.0.0/24
        [0]    directly connected, GigabitEthernet0/6/0/0
    10.0.0.0/8    auto-summary
    12.0.0.0/24
        [5] distance: 20    redistributed
    12.0.0.0/8    auto-summary
    50.50.0.0/24
        [1] via 10.0.0.20, next-hop 10.0.0.20, Uptime: 1s, GigabitEthernet0/6/0/0
    50.50.1.0/24 (inactive)
        [1] via 10.0.0.20, next-hop 10.0.0.20, Uptime: 1s, GigabitEthernet0/6/0/0
    50.0.0.0/8    auto-summary
    90.90.0.0/24
        [5] distance: 20    redistributed
    90.90.1.0/24
        [5] distance: 20    redistributed
```

The following is sample output from the **show rip ipv6 database** command:

```
Router# show rip ipv6 database
Mon Mar 22 21:07:22.465 IST

IPv6 Routes held in RIP's topology database:
2001:DB8::2/48
    [0]    directly connected, GigabitEthernet0/0/0/0
2001:DB8::1/48
    [0]    directly connected, Loopback0
```

This table describes the significant fields shown in the display.

Table 169: show rip IPv4 database Field Descriptions

Field	Description
10.0.0.0/24 [0] directly connected, GigabitEthernet0/6/0/0	Prefix and prefix length for a RIP connected route. 10.0.0.0/24 is directly connected to GigabitEthernet 0/6/0/0. The [0] represents the metric.
10.0.0.0/8 auto-summary	10.0.0.0/8 is a summary route entry.
12.0.0.0/24 [5] distance: 20 redistributed	12.0.0.0/24 is a redistributed route. The metric is 5, and the distance is 20.

Field	Description
50.50.0.0/24 [1] via 10.0.0.20, next-hop 10.0.0.20, Uptime: 1s, GigabitEthernet0/6/0/0	The destination route 50.50.0.0/24 is learned through RIP, and the source 10.0.0.20 advertised it from GigabitEthernet 0/6/0/0. The route was last updated one second ago.
50.50.1.0/24 (inactive) [1] via 10.0.0.20, next hop 10.0.0.20, Uptime: 1s, GigabitEthernet0/6/0/0	The destination route 50.50.1.0/24 is not active in the routing table.

This table describes the significant fields shown in the display.

Table 170: show rip IPv6 database Field Descriptions

Field	Description
2001:DB8::2/48 [0] directly connected, GigabitEthernet0/6/0/0	Prefix and prefix length for a RIP connected route. 2001:DB8::2/48 is directly connected to GigabitEthernet 0/0/0/0. The [0] represents the metric.
2001:DB8::1/48 [0] directly connected, Loopback0	2001:DB8::1/48 is directly connected to Loopback 0/0/0/0. The [0] represents the metric. .

show rip interface

To display interface entry information from the Routing Information Protocol (RIP) topology table, use the **show rip interface**command in EXEC mode.

show rip [**vrf** {*vrf-name* | **all**}] **interface** [*type interface-path-id*]

Syntax Description

vrf { *vrf* | **all** }	(Optional) Specifies a particular VPN routing and forwarding (VRF) instance or all VRF instances.
interface	(Optional) Specifies the interface from which to clear topology entries.
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

No default behavior or values

Command Modes

EXEC

Command History

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
rip	read

Examples

This example is sample output from the **show rip interface** command:

```
RP/0/RSP0/CPU0:router# show rip interface

GigabitEthernet0_6_0_0
Rip enabled?:               Yes
Out-of-memory state:        Normal
Broadcast for V2:           No
Accept Metric 0?:           No
Send versions:              2
Receive versions:           2
Interface state:            Up
IP address:                 10.0.0.12/24
Metric Cost:                0
```

```
Split horizon:              Enabled
Poison Reverse:             Disabled
Joined multicast group?:    Yes

GigabitEthernet0_6_0_2
Out-of-memory state:        Normal
Rip enabled?:               Yes
Broadcast for V2:           No
Accept Metric 0?:           No
Send versions:              2
Receive versions:           2
Interface state:            Up
IP address:                 12.0.0.12/24
Metric Cost:                0
Split horizon:              Enabled
Poison Reverse:             Disabled
Joined multicast group?:    Yes

RIP peers attached to this interface:
    12.0.0.13
        uptime: 3    version: 2
        packets discarded: 0    routes discarded: 402
```

This table describes the significant fields shown in the display.

Table 171: show rip interface Field Descriptions

Field	Description
Rip enabled?	Specifies whether the RIP routing protocol is enabled on the interface.
Out-of-memory state	Specifies the current out-of-memory state on the interface.
Broadcast for V2	Specifies whether RIP Version 2 output packets are sent to a broadcast address on the interface.
Accept Metric 0?	Specifies whether this interface accepts routing entries from RIP updates with a metric set to zero (0).
Send versions:	Specifies which version RIP uses to send out packets on this interface.
Receive versions:	Specifies which version packets RIP accepts on this interface.
Interface state:	Specifies whether the interface is in an up or a down state.
IP address:	IP address of the interface.
Metric Cost:	Specifies metric cost value.
Split horizon:	Specifies whether split horizon is enabled on this interface.
Poison Reverse:	Specifies whether poison reverse is enabled on this interface.
Joined multicast group?:	Specifies whether the interface has joined the RIP multicast group 224.0.0.9.

Field	Description
RIP peers attached to this interface 12.0.0.13	List of RIP neighbors on this interface.
uptime: 3	Specifies how long this neighbor is up.
version: 2	Specifies which version packets are received from this neighbor.
packets discarded: 0	Specifies the number of packets discarded from this neighbor.
routes discarded: 402	Specifies the number of routes discarded from this neighbor.

show rip statistics

To display statistical entry information from the Routing Information Protocol (RIP) topology table, use the **show rip statistics** command in EXEC mode.

show rip { **ipv4** | **ipv6** } [**vrf** { *vrf-name* | **all** }] **statistics**

Syntax Description

ipv4 | **ipv6**	Specifiies IPv4 or IPv6.
vrf { *vrf* | **all** }	(Optional) Specifies a particular VPN routing and forwarding (VRF) instance or all VRF instances.

Command Default

No default behavior or values

Command Modes

EXEC

Command History

Release	Modification
Release 3.7.2	This command was introduced.
Release 7.5.2	Support for IPv6 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
rip	read

Examples

The following example is sample output from the **show rip ipv4 statistics** command:

```
Router# show rip ipv4 statistics

RIP statistics:
    Total messages sent:        5597
    Message send failures:      0
    Regular updates sent:       5566
    Queries responsed to:       0
    RIB updates:                6
    Total packets received:     5743
    Discarded packets:          0
    Discarded routes:           0
    Number of routes allocated: 18
    Number of paths allocated:  14

    Route malloc failures:      0
    Path malloc failures:       0
```

The following example is sample output from the **show rip ipv6 statistics** command:

```
  Router# show rip ipv6 statistics

Mon Mar 22 21:10:30.110 IST

RIP statistics:
Total messages sent:        79
Message send failures:      0
Regular updates sent:       76
Queries responsed to:       1
RIB updates:                2
Total packets received:     38
Discarded packets:          0
Discarded routes:           0
Packet received at standby: 0
Number of routes allocated: 13
Number of paths allocated:  7
Route malloc failures:      0
Path malloc failures:       0
```

This table describes the significant fields shown in the display.

Table 172: show rip statistics Field Descriptions

Field	Description
Total messages sent	Number of RIP packets sent.
Message send failures	Number of times that the packet send operation failed.
Queries responsed to	Number of times RIP updates are sent in response to a RIP query.
RIB updates	Number of route addition and deletion messages sent to RIB.
Total packets received	Number of RIP packets received.
Discarded packets	Number of received RIP packets that are discarded.
Discarded routes	Number of routes (in received RIP update packets) that are discarded.
Number of routes allocated	Number of routes allocated for the RIP internal topology database.
Number of paths allocated	Number of paths allocated for the RIP internal topology database.
Route malloc failures	Number of failures during route allocation.
Path malloc failures	Number of failures during route allocation.

Note The number of routes found in the allocated field might ot be the same number of routes present in the RIP database.

site-of-origin (RIP)

To configure the Site of Origin (SoO) filtering on a Routing Information Protocol (RIP) interface, use the **site-of-origin** command in interface configuration mode. To disable SoO filtering on an interface, use the **no** form of this command.

site-of-origin {*as-number* **:** *number* | *ip-address* **:** *number*}
no site-of-origin {*as-number* **:** *number* | *ip-address* **:** *number*}

Syntax Description

as-number :	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. A colon is used to separate the autonomous system number and network number.
number	Network number. Range is from 0 to 4294967295 when a 2-byte AS number is used. Range is from 0 to 65535 when a 4-byte AS number is used.
ip-address	The IP address argument specifies the IP address in four-part, dotted-decimal notation. A colon is used to separate the IP address and network number.

Command Default

No default behavior or values

Command Modes

Interface configuration

Command History

Release	Modification
Release 3.7.2	This command was introduced.
Release 3.9.0	Asplain format for 4-byte Autonomous system numbers notation was supported.

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.

A RIP process must be capable of retrieving the SoO attribute on routes redistributed from the Border Gateway Protocol (BGP) when required to support complex topologies that include MPLS VPN links between sites with backdoor links.

Use the **site-of-origin** command to set an SoO BGP extended community attribute that is used to identify routes that have originated from a site so that the readvertisement of that prefix back to the source site can be prevented. The SoO extended community uniquely identifies the site from which a provider edge (PE) router has learned a route.

Task ID

Task ID	Operations
rip	read, write

Examples

The following example shows how to configure SoO filtering on a RIP interface:

```
RP/0/RSP0/CPU0:router(config)# router rip
RP/0/RSP0/CPU0:router(config-rip) interface GigabitEthernet 0/1/0/0
RP/0/RSP0/CPU0:router(config-rip-if)# site-of-origin 10.0.0.1:20
```

split-horizon disable (RIP)

To disable split horizon for a Routing Information Protocol (RIP) process, use the **split-horizon disable** command in interface configuration mode. To enable split horizon, use the **no** form of this command.

split-horizon disable
no split-horizon disable

Syntax Description

This command has no arguments or keywords.

Command Default

Split horizon is enabled for a RIP process.

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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

You can explicitly specify the **split-horizon disable** command in your configuration.

If split horizon is disabled, the poison reverse configuration is ignored.



Note In general, we recommend that you do not change the default state of split horizon unless you are certain that your application requires the change to properly advertise routes.

Task ID

Task ID	Operations
rip	read, write

Examples

The following example shows how to disable split horizon on a Packet-over-SONET/SDH link:

```
Router(config)# router rip
Router(config-rip)# interface GigabitEthernet 0/1/0/0
Router(config-rip-if)# split-horizon disable
```

Related Commands

Command	Description
poison-reverse, on page 1561	Enables poison reverse processing of RIP router updates.

timers basic

To adjust Routing Information Protocol (RIP) network timers, use the **timers basic** command in router configuration mode. To restore the timers default values, use the **no** form of this command.

timers basic *update invalid holddown flush*
no timers basic

Syntax Description

update	Rate, in seconds, at which updates are sent. This is the fundamental timing parameter of the routing protocol. Range is 5 to 50000.
invalid	Interval, in seconds, after which a route is declared invalid; it should be at least three times the value of the update argument. A route becomes invalid when there is an absence of updates that refresh the route. The route then enters into a holddown state. The route is marked inaccessible and is advertised as unreachable. Range is 15 to 200000.
holddown	Interval, in seconds, during which routing information regarding better paths is suppressed. It should be at least three times the value of the update argument. A route enters into a holddown state when an update packet is received that indicates that the route is unreachable. The route is marked inaccessible and is advertised as unreachable. When holddown expires, routes advertised by other sources are accepted, and the route is no longer inaccessible. Range is 15 to 200000.
flush	Amount of time, in seconds, that must pass before the route is removed from the routing table; the interval specified should be greater than the value of the *invalid* argument. If it is less than the invalid timer value, the proper holddown interval cannot elapse, which results in a new route being accepted before the holddown interval expires. Range is 16 to 250000.

Command Default

update : 30

invalid : 180

holddown : 180

flush : 240

Command Modes

Router 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

The basic timing parameters for RIP are adjustable. Because RIP is running a distributed, asynchronous routing algorithm, these timers must be the same for all routers in the network.

Note Use the **show rip** command to display the current and default timer values.

Task ID

Task ID	Operations
rip	read, write

Examples

The following example shows how to set updates to be broadcast every 5 seconds. If a router is not heard from in 15 seconds, the route is declared unusable. Further information is suppressed for an additional 15 seconds. At the end of the flush period, the route is flushed from the routing table.

```
RP/0/RSP0/CPU0:router(config)# router rip
RP/0/RSP0/CPU0:router(config-rip) timers basic 5 15 15 30
```

Related Commands

Command	Description
show rip, on page 1575	Displays configuration and status of RIP.

validate-update-source disable

To stop the Cisco IOS XR software from validating the source IP address of incoming routing updates for Routing Information Protocol (RIP), use the **validate-update-source disable** command in router configuration mode. To reenable this function, use the **no** form of this command.

validate-update-source disable
no validate-update-source disable

Syntax Description

This command has no arguments or keywords.

Command Default

The source IP address of incoming updates for RIP is always validated.

Command Modes

Router 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

When the **validate-update-source disable** command is used, validation is not performed.

By default, the software ensures that the source IP address of incoming routing updates is on the same IP network as one of the addresses defined for the receiving interface.

For unnumbered IP interfaces (interfaces configured as IP unnumbered), no checking is performed.

Task ID

Task ID	Operations
rip	read, write

Examples

The following example shows how to disable source validation:

```
RP/0/RSP0/CPU0:router(config)# router rip
RP/0/RSP0/CPU0:router(config-rip) validate-update-source disable
```

vrf (RIP)

To define a VPN routing and forwarding (VRF) instance and to enter VRF configuration mode, use the **vrf** command in router configuration mode. To remove a VRF instance use the no form of this command.

vrf *vrf-name*
no vrf *vrf-name*

Syntax Description

vrf-name	Specifies a particular VPN routing and forwarding instance.

Command Default

No VRFs are defined.

Command Modes

Router 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

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.

From VRF configuration mode, you can issue all commands available in router configuration mode such as the **auto-summary** command.

Task ID

Task ID	Operations
rip	read, write

Examples

The following example shows how to enter VRF configuration mode and identify RIP commands that can be issued from that mode:

```
RP/0/RSP0/CPU0:router(config)# router rip
RP/0/RSP0/CPU0:router(config-rip)# vrf vpn-1
RP/0/RSP0/CPU0:router(config-rip-vrf)# ?

   auto-summary            Enable automatic network number summarization
   broadcast-for-v2        Send RIP v2 output packets to broadcast address
   commit                  Commit the configuration changes to running
   default-information     Control distribution of default information
   default-metric          Set metric of redistributed routes
   describe                Describe a command without taking real actions
   distance                Define an administrative distance
   do                      Run an exec command
   exit                    Exit from this submode
```

```
   interface               Enter the RIP interface configuration submode
   maximum-paths           Maximum number of paths allowed per route
   neighbor                Specify a neighbor router
   no                      Negate a command or set its defaults
   nsf                     Enable Cisco Non Stop Forwarding
   output-delay            Interpacket delay for RIP updates
   redistribute            Redistribute information from another routing protocol
   route-policy            Apply route policy to routing updates
   show                    Show contents of configuration
   timers                  Adjust routing timers
   validate-update-source  Validate source address of routing updates
RP/0/RSP0/CPU0:router(config-rip-vrf)#
```



Routing Policy Language Commands

This module describes the Cisco IOS XR software routing policy language (RPL) commands used to create, modify, monitor, and maintain routing policies.

For detailed information about RPL concepts, configuration tasks, and examples, see the *Implementing Routing Policy on* Cisco ASR 9000 Series Router module in the *Routing Configuration Guide for Cisco ASR 9000 Series Routers* .

- abort (RPL), on page 1599
- add, on page 1601
- apply, on page 1603
- as-path in, on page 1605
- as-path is-local, on page 1607
- as-path length, on page 1609
- as-path neighbor-is, on page 1611
- as-path originates-from, on page 1614
- as-path passes-through, on page 1616
- as-set, on page 1618
- as-path-set, on page 1620
- as-path unique-length, on page 1622
- community is-empty, on page 1624
- community matches-any, on page 1625
- community matches-every, on page 1627
- community matches-within, on page 1629
- community-set, on page 1631
- delete community, on page 1634
- delete large-community, on page 1636
- delete extcommunity rt, on page 1638
- destination in, on page 1640
- done, on page 1642
- drop, on page 1644
- edit, on page 1646
- end-global, on page 1649
- end-policy, on page 1650
- end-set, on page 1651
- event prefix route-policy pass, on page 1653

- extcommunity rt is-empty, on page 1654
- extcommunity rt matches-any, on page 1655
- extcommunity rt matches-every, on page 1657
- extcommunity rt matches-within, on page 1659
- extcommunity-set bandwidth, on page 1661
- extcommunity-set cost, on page 1662
- extcommunity-set opaque, on page 1664
- extcommunity-set rt, on page 1665
- extcommunity-set seg-nh, on page 1667
- extcommunity-set soo, on page 1668
- extcommunity soo is-empty, on page 1670
- extcommunity soo matches-any, on page 1671
- extcommunity soo matches-every, on page 1673
- globalVar*N* is, on page 1675
- if, on page 1677
- **if route-aggregated** , on page 1683
- is-best-path, on page 1684
- is-backup-path, on page 1685
- is-multi-path, on page 1686
- local-preference, on page 1687
- large-community is-empty, on page 1688
- large-community matches-any, on page 1690
- large-community matches-every, on page 1692
- large-community matches-within, on page 1694
- large-community-set, on page 1696
- med, on page 1697
- next-hop in, on page 1699
- orf prefix in, on page 1701
- origin is, on page 1703
- ospf-area, on page 1705
- ospf-area-set, on page 1707
- pass, on page 1709
- path-type is, on page 1711
- policy-global, on page 1712
- prefix-set, on page 1714
- prepend as-path, on page 1716
- protocol, on page 1718
- rd in, on page 1720
- rd-set, on page 1721
- replace as-path, on page 1723
- remove as-path private-as, on page 1725
- rib-has-route, on page 1726
- route-has-label, on page 1728
- route-policy (RPL), on page 1729
- route-type is, on page 1731
- rpl editor, on page 1733

- rpl maximum, on page 1735
- rpl set-exit-as-abort, on page 1737
- set administrative-distance, on page 1738
- set aigp-metric, on page 1739
- set community, on page 1740
- set core-tree, on page 1742
- set dampening, on page 1743
- set fallback-vrf-lookup, on page 1745
- set eigrp-metric, on page 1746
- set extcommunity color, on page 1748
- set extcommunity cost, on page 1749
- set extcommunity rt, on page 1751
- set ip-precedence, on page 1753
- set isis-metric, on page 1755
- set label, on page 1756
- set label-mode, on page 1757
- set large-community, on page 1759
- set level, on page 1761
- set local-preference, on page 1762
- set med, on page 1763
- set metric-type (IS-IS), on page 1765
- set metric-type (OSPF), on page 1766
- set next-hop, on page 1767
- set origin, on page 1769
- set ospf-metric, on page 1770
- set path-selection, on page 1771
- set qos-group (RPL), on page 1773
- set rib-metric, on page 1774
- set rip-metric, on page 1775
- set rip-tag, on page 1776
- set rpf-topology, on page 1777
- set rtset route-limit , on page 1779
- set spf-priority, on page 1780
- set tag, on page 1781
- set traffic-index, on page 1782
- set vpn-distinguisher, on page 1784
- set weight, on page 1785
- show rpl, on page 1787
- show rpl active as-path-set, on page 1789
- show rpl active community-set, on page 1792
- show rpl active extcommunity-set, on page 1795
- show rpl active prefix-set, on page 1798
- show rpl active rd-set, on page 1801
- show rpl active route-policy, on page 1803
- show rpl as-path-set, on page 1805
- show rpl as-path-set attachpoints, on page 1807

- show rpl as-path-set references, on page 1810
- show rpl community-set, on page 1812
- show rpl community-set attachpoints, on page 1814
- show rpl community-set references, on page 1816
- show rpl extcommunity-set, on page 1819
- show rpl inactive as-path-set, on page 1822
- show rpl inactive community-set, on page 1825
- show rpl inactive extcommunity-set, on page 1828
- show rpl inactive prefix-set, on page 1831
- show rpl inactive rd-set, on page 1834
- show rpl inactive route-policy, on page 1836
- show rpl maximum, on page 1839
- show rpl policy-global references, on page 1841
- show rpl prefix-set, on page 1843
- show rpl prefix-set attachpoints, on page 1845
- show rpl prefix-set references, on page 1848
- show rpl rd-set, on page 1850
- show rpl rd-set attachpoints, on page 1852
- show rpl rd-set references, on page 1854
- show rpl route-policy, on page 1856
- show rpl route-policy attachpoints, on page 1859
- show rpl route-policy inline, on page 1862
- show rpl route-policy references, on page 1864
- show rpl route-policy uses, on page 1867
- show rpl unused as-path-set, on page 1870
- show rpl unused community-set, on page 1873
- show rpl unused extcommunity-set, on page 1876
- show rpl unused prefix-set, on page 1878
- show rpl unused rd-set, on page 1881
- show rpl unused route-policy, on page 1883
- source in, on page 1886
- source rt, on page 1888
- suppress-route, on page 1889
- tag, on page 1890
- tag in, on page 1891
- tag-set, on page 1893
- unsuppress-route, on page 1894
- var globalVar*N*, on page 1896
- vpn-distinguisher is, on page 1897

abort (RPL)

To discard a route policy or set definition and return to global configuration mode, use the **abort** command in the appropriate configuration mode.

abort

Syntax Description

This command has no keywords or arguments.

This command has no arguments or keywords.

Command Default

No default behavior or values

Command Modes

Route-policy configuration

Prefix set configuration

Route distinguisher set configuration

AS path set configuration

Community set configuration

Extended community set 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 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
route-policy	read, write

Examples

The following example shows how to discard a route policy definition that was started and return to global configuration mode:

```
RP/0/RSP0/CPU0:router(config)# route-policy policy_1
RP/0/RSP0/CPU0:router(config-rpl)# if as-path is-local then
RP/0/RSP0/CPU0:router(config-rpl-if)# abort
RP/0/RSP0/CPU0:router(config)#
```

The following example shows how to discard a prefix set definition that was started and return to global configuration mode:

```
RP/0/RSP0/CPU0:router(config)# prefix-set legal-ipv4-prefix-examples
RP/0/RSP0/CPU0:router(config-pfx)# 10.0.1.1,
RP/0/RSP0/CPU0:router(config-pfx)# 10.0.2.0/24,
RP/0/RSP0/CPU0:router(config-pfx)# abort
RP/0/RSP0/CPU0:router(config)#
```

add

To add a value to an Routing Information Protocol (RIP) or Enhanced Interior Gateway Protocol (EIGRP) existing metric, use the **add** command in route-policy configuration mode.

add {**eigrp-metric** *bandwidth delay reliability loading max-transmission* | **rip-metric** {*numberparameter*}}

Syntax Description

eigrp-metric	Specifies an EIGRP metric attribute.
bandwidth	Bandwidth in kilobits per second. Range is from 0 to 4294967295.
delay	Delay in 10-microsecond units. Range is from 0 to 4294967295.
reliability	Reliability metric. 255 is 100-percent reliable. Range is from 0 to 255.
loading	Effective bandwidth (loading). 255 is 100-percent loaded. Range is from 0 to 255.
max-transmission	Maximum transmission of the path. Range is from 0 to 65535.
rip-metric	Specifies an RIP metric attribute.
number	Value assigned to a four-bit unsigned integer. Range is from 0 to 16.
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 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

If the add value is greater than the maximum allowed value, the metric is added. If the resulting metric exceeds the maximum for the routing protocol, then the route is dropped (by the client routing protocol).

Task ID

Task ID	Operations
route-policy	read, write

Examples

The following example shows how to offset the RIP metric value:

```
RP/0/RSP0/CPU0:router(config)# route-policy policy_1
```

```
RP/0/RSP0/CPU0:router(config-rpl)# add rip-metric 4
RP/0/RSP0/CPU0:router(config-rpl)# end-policy
```

The following example shows how to set the EIGRP metric value:

```
RP/0/RSP0/CPU0:router(config)# route-policy policy_1
RP/0/RSP0/CPU0:router(config-rpl)# add eigrp-metric 50000 24000 230 14000
RP/0/RSP0/CPU0:router(config-rpl)# end-policy
```

apply

To execute a parameterized or unparameterized policy from within another policy, use the **apply** command in route-policy configuration mode.

apply *policy_name* [*argument1, argument2, . . . , argumentN*]

Syntax Description

policy_name	Name of a route policy.
argument	(Optional) Parameter name. The *argument* can be a value (for example, '100') or a parameter (for example, '$parameter')

Command Default

No default behavior or values

Command Modes

Route-policy 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **apply** command to execute a policy (either parameterized or unparameterized) from within another policy, which allows for the reuse of common blocks of policy.

Task ID

Task ID	Operations
route-policy	read, write

Examples

In the following example, the policy CustomerIn applies the route-policy SetLocalPref to conditionally set the local preference on a route. The parameters 20, 30, 40, and 50 are passed to the parameterized policy SetLocalPref, where the local preference is set to:

- 20, if the community 217:20 is present in the route
- 30, if the community 217:30 is present in the route
- 40, if the community 217:40 is present in the route
- 50, if the community 217:50 is present in the route

```
RP/0/RSP0/CPU0:router(config)# route-policy SetLocalPref ($lp0, $lp1, $lp2, $lp3, $lp4)
RP/0/RSP0/CPU0:router(config-rpl)# if community matches-any ($lp0:$lp1)then
RP/0/RSP0/CPU0:router(config-rpl-elseif)# set local-preference $lp1
RP/0/RSP0/CPU0:router(config-rpl-elseif)# elseif community matches-any ($lp0:$lp2) then
RP/0/RSP0/CPU0:router(config-rpl-elseif)# set local-preference $lp2
RP/0/RSP0/CPU0:router(config-rpl-elseif)# elseif community matches-any ($lp0:$lp3) then
RP/0/RSP0/CPU0:router(config-rpl-elseif)# set local-preference $lp3
RP/0/RSP0/CPU0:router(config-rpl-elseif)# elseif community matches-any ($lp0:$lp4) then
```

```
RP/0/RSP0/CPU0:router(config-rpl-elseif)# set local-preference $lp4
RP/0/RSP0/CPU0:router(config-rpl-elseif)# endif
RP/0/RSP0/CPU0:router(config-rpl)# end-policy

RP/0/RSP0/CPU0:router(config)# route-policy CustomerIn($cust)
RP/0/RSP0/CPU0:router(config-rpl)# apply SetLocalPref ($cust, 20, 30, 40, 50)
RP/0/RSP0/CPU0:router(config-rpl)# end-policy

RP/0/RSP0/CPU0:router(config)# route-policy Cust_217
RP/0/RSP0/CPU0:router(config-rpl)# apply CustomerIn(217)
RP/0/RSP0/CPU0:router(config-rpl)# end-policy
```

as-path in

To match the AS path of a route to an AS path set, use the **as-path in** command in route-policy configuration mode.

as-path in {*as-path-set-nameinline-as-path-setparameter*}

Syntax Description

as-path-set-name	Name of an AS path set.
inline-as-path-set	Inline AS path set. The inline AS path set must be enclosed in parentheses.
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 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **as-path in** command as a conditional expression within an **if** statement to match the AS path of a route to an AS path set. The AS path is a sequence of autonomous system numbers traversed by a route.

Note For a list of all conditional expressions available within an **if** statement, see the **if** command.

The **as-path in** command evaluates to true if at least one of the regular expressions defined in the associated AS path set matches the AS path attribute of the route.

In the case where the AS path set is defined but contains no elements in it, the **as-path in** conditional expression command returns false.

Task ID

Task ID	Operations
route-policy	read, write

Examples

For example, assume we have an AS path set named my-as-set defined as follows:

```
RP/0/RSP0/CPU0:router(config)# as-path-set my-as-set
RP/0/RSP0/CPU0:router(config-as)# ios-regex '_12$',
RP/0/RSP0/CPU0:router(config-as)# ios-regex '_13$'
```

```
RP/0/RSP0/CPU0:router(config-as)# end-set
```

and the following policy excerpt using an *as-path-set-name* argument:

```
RP/0/RSP0/CPU0:router(config-rpl)# if as-path in my-as-set then
RP/0/RSP0/CPU0:router(config-rpl-if)# set local-preference 100
RP/0/RSP0/CPU0:router(config-rpl-if)# endif
RP/0/RSP0/CPU0:router(config-rpl)#
```

The AS path in condition evaluates to true if one or more of the regular expression matches in the set my-as-set match the AS path associated with the route. In the case of a defined but empty AS path set, this operator returns false.

The preceding policy excerpt is equivalent to the following version, which uses an *inline-as-path* set variable:

```
RP/0/RSP0/CPU0:router(config-rpl)# if as-path in (ios-regex '_12$,ios-regex '_13$') then
RP/0/RSP0/CPU0:router(config-rpl-if)# set local-preference 100
RP/0/RSP0/CPU0:router(config-rpl-if)# endif
RP/0/RSP0/CPU0:router(config-rpl)#
```

as-path is-local

To determine if this router or another router within this autonomous system or confederation originated a Border Gateway Protocol (BGP) route, use the **as-path is-local** command in route-policy configuration mode.

as-path is-local

Syntax Description

This command has no arguments or keywords.

Command Default

No default behavior or values

Command Modes

Route-policy 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **as-path is-local** command as a conditional expression within an **if** statement to determine if this router (or another router within this autonomous system or confederation) originated the route.

Note For a list of all conditional expressions available within an **if** statement, see the **if** command.

Routes that are locally originated within the autonomous system or confederation carry an empty AS path. For the Border Gateway Protocol (BGP) specification, when a route is advertised across the autonomous system boundary or a confederation boundary, the local autonomous system number or confederation ID is appended to the autonomous system path. The AS path of a locally originated aggregate is also empty unless it has been modified by policy.

The **is-local** operator evaluates to true for autonomous system paths that are empty. An empty AS path is how an AS path that is local to our autonomous system is represented in BGP.

Task ID

Task ID	Operations
route-policy	read, write

Examples

In the following example, if the AS path is local, then the local preference is set to 100:

```
RP/0/RSP0/CPU0:router(config-rpl)# if as-path is-local then
RP/0/RSP0/CPU0:router(config-rpl-if)# set local-preference 100
RP/0/RSP0/CPU0:router(config-rpl-if)# endif
```

```
RP/0/RSP0/CPU0:router(config-rpl)#
```

as-path length

To compare the number of ASN in the AS path of a Border Gateway Protocol (BGP) route, use the **as-path length** command in route-policy configuration mode.

as-path length {**eq** | **is** | **ge** | **le**} {*numberparameter*}

Syntax Description

eq | **is** | **ge** | **le**	Equal to; greater than or equal to; less than or equal to.
number	Value assigned to an 11-bit unsigned integer. Range is from 0 to 2047.
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 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **as-path length** command as a conditional expression within an **if** statement to perform a conditional check based on the length of the AS path.

Note For a list of all conditional expressions available within an **if** statement, see the **if** command.

This command takes either a specific integer value or a range of integer values specified with the **ge** and **le** operators. Any or all these integers can be parameterized. The operator counts one for each autonomous system in the path. In the case where the route may be aggregated and contain one or more AS sets, the length operator adds one for each set present, the occurrence of an AS set typically indicates that this route is an aggregated route, and the aggregated route had a component route that contained one of the autonomous systems in the set. Likewise, in the case of confederations, a count of one is added for each confederation in the path or each confederation set in the path. A null AS path has a length of zero.

Task ID

Task ID	Operations
route-policy	read, write

Examples

In the following example, if the AS path length equals 10, then the local preference is set to 100:

```
RP/0/RSP0/CPU0:router(config-rpl)# if as-path length eq 10 then
```

```
RP/0/RSP0/CPU0:router(config-rpl-if)# set local-preference 100
RP/0/RSP0/CPU0:router(config-rpl-if)# endif
RP/0/RSP0/CPU0:router(config-rpl)#
```

Related Commands

Command	Description
as-path in, on page 1605	Matches the AS path of a route to an AS path set.
as-path originates-from, on page 1614	Compares an AS path against the AS sequence beginning with the AS number that originated a route
as-path passes-through, on page 1616	Verifies if the supplied integer or parameter appears anywhere in the AS path or if the supplied sequence of integers and parameters appears, in the same order, anywhere in the AS path.
as-path unique-length, on page 1622	Performs specific checks based on the length of the AS path.

as-path neighbor-is

To test autonomous system numbers at the head of the AS path against a sequence of one or more values or parameters, use the **as-path neighbor-is** command in route-policy configuration mode.

as-path neighbor-is *as-number-list* [**exact**]

Syntax Description

as-number-list	Numbers or parameters, enclosed in single quotation marks, that represent a sequence of autonomous system numbers. • 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.
exact	(Optional) Specifies that with the **exact** keyword, the *as-number-list* value must identically match the AS path for the route; without the **exact** keyword, any element in the *as-number-list* argument matches one or more occurrences of that element in the AS path for the route.

Command Default

No default behavior or values

Command Modes

Route-policy configuration

Command History

Release	Modification
Release 3.7.2	This command was introduced.
Release 3.9.0	Asplain format for 4-byte Autonomous system numbers notation was supported.

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.

Use the **as-path neighbor-is** command as a conditional expression within an **if** statement to test the autonomous system number or numbers at the head of the AS path against a sequence of one or more integral values or parameters. In other words, to test to learn if the sequence of autonomous system numbers matches the path beginning with the neighboring autonomous system from which this route was heard.



Note For a list of all conditional expressions available within an **if** statement, see the **if** command.

This command has an equivalent regular expression (ios-regex). For example, AS path neighbor-is '**1**' would be '**^1_**'.

Task ID

Task ID	Operations
route-policy	read, write

Examples

The following are incomplete configuration examples:

```
RP/0/RSP0/CPU0:router(config-rpl)# if as-path neighbor-is '10' then
RP/0/RSP0/CPU0:router(config-rpl-if)# if as-path neighbor-is '$asnum' then
RP/0/RSP0/CPU0:router(config-rpl-if)# if as-path neighbor-is '10 20' then
```

These statements evaluate to true when the first autonomous system numbers on the AS path match, in the same order, the supplied parameters or integer values in the **neighbor-is** statement. If the neighboring autonomous system location happens to be an AS-set, the operator evaluates to true if the corresponding argument to the **neighbor-is** operator is an element of the AS-set.

Without the **exact** keyword, repeated autonomous system numbers in the AS path are ignored. For example,

```
RP/0/RSP0/CPU0:router(config-rpl)# if as-path neighbor-is '10 20' then
```

matches an AS path beginning

```
10 10 10 20 ...
```

and an AS path beginning:

```
10 20 ....
```

With the **exact** keyword, repetitions are not ignored, therefore

```
RP/0/RSP0/CPU0:router(config-rpl)# if as-path neighbor-is '10 20' exact then
```

matches the second of these AS paths but not the first.

Related Commands

Command	Description
as-path in, on page 1605	Matches the AS path of a route to an AS path set.
as-path length, on page 1609	Compares the number of ASN in the AS path of a route.
as-path originates-from, on page 1614	Compares an AS path to the AS sequence beginning with the AS number that originated a route.
as-path passes-through, on page 1616	Verifies if the supplied integer or parameter appears anywhere in the AS path or if the supplied sequence of integers and parameters appears, in the same order, anywhere in the AS path.

Command	Description
as-path unique-length, on page 1622	Performs specific checks based on the length of the AS path.

as-path originates-from

To compare an AS path against the AS sequence beginning with the AS number that originated a route, use the **as-path originates-from** command in route-policy configuration mode.

as-path originates-from *as-number-list* [**exact**]

Syntax Description

as-number-list	Numbers or parameters, enclosed in single quotation marks, that represent a sequence of autonomous system numbers. • 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.
exact	(Optional) Specifies that with the **exact** keyword, the *as-number-list* value must identically match the AS path for the route; without the **exact** keyword, any element in the *as-number-list* argument matches one or more occurrences of that element in the AS path for the route.

Command Default

No default behavior or values

Command Modes

Route-policy configuration

Command History

Release	Modification
Release 3.7.2	This command was introduced.
Release 3.9.0	Asplain format for 4-byte Autonomous system numbers notation was supported.

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.

Use the **as-path originates-from** command as a conditional expression within an **if** statement to compare an AS path to the autonomous system sequence.



Note For a list of all conditional expressions available within an **if** statement, see the **if** command.

The **originates-from** operator is similar to the **neighbor-is** operator, except that it looks at the autonomous system number at the opposite end of the AS path. In other words, it is comparing to the autonomous system that originated the route. It can take numbers or parameters, enclosed in single quotation marks, that represent a sequence of autonomous system numbers. When more than one number is specified in the list, the sequence of autonomous system numbers listed must appear as a subsequence in the AS path, with the last number corresponding to the autonomous system that originated the route.

Task ID

Task ID	Operations
route-policy	read, write

Examples

The following are incomplete configuration examples:

```
RP/0/RSP0/CPU0:router(config-rpl)# if as-path originates-from '10 11' then
RP/0/RSP0/CPU0:router(config-rpl-if)# if as-path originates-from '$asnum 11' then
```

The first line of the preceding example evaluates to true if autonomous system 11 originated the route and then advertised it to autonomous system 10, from which the route was eventually propagated to us. In the case where the route has been aggregated, and the location of the originating autonomous system contains an AS-set, the **originates-from** operator evaluates to true if the argument to the **originates-from** operator is contained in the AS-set.

Without the **exact** keyword, repeated autonomous system numbers in the AS path are ignored. For example,

```
RP/0/RSP0/CPU0:router(config-rpl)# if as-path originates-from '10 11' then
```

matches an autonomous system path ending

```
...10 10 10 11
```

and an autonomous system path ending

```
...10 11
```

With the **exact** keyword, repetitions are not ignored, therefore

```
RP/0/RSP0/CPU0:router(config-rpl)# if as-path originates-from '10 11' exact then
```

matches the second of these autonomous system paths but not the first.

Related Commands

Command	Description
as-path in, on page 1605	Matches the AS path of a route to an AS path set.
as-path length, on page 1609	Compares the number of ASN in the AS path of a route
as-path passes-through, on page 1616	Verifies if the supplied integer or parameter appears anywhere in the AS path or if the supplied sequence of integers and parameters appears, in the same order
as-path unique-length, on page 1622	Performs specific checks based on the length of the AS path.

as-path passes-through

To verify if the supplied integer or parameter appears anywhere in the AS path or if the supplied sequence of integers and parameters appears, in the same order, anywhere in the AS path, use the **as-path passes-through** command in route-policy configuration mode.

as-path passes-through *as-number-list* [**exact**]

Syntax Description

as-number-list	Numbers or parameters, enclosed in single quotation marks, that represent a sequence of autonomous system numbers. • 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.
exact	(Optional) Specifies that with the **exact** keyword, the *as-number-list* value must identically match the AS path for the route; without the **exact** keyword, any element in the *as-number-list* argument matches one or more occurrences of that element in the AS path for the route.

Command Default

No default behavior or values

Command Modes

Route-policy configuration

Command History

Release	Modification
Release 3.7.2	This command was introduced.
Release 3.9.0	Asplain format for 4-byte Autonomous system numbers notation was supported.

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.

Use the **as-path passes-through** command as a conditional expression within an **if** statement to verify if the specified integer or parameter appears anywhere in the AS path or if the sequence of integers and parameters appears.

Note For a list of all conditional expressions available within an **if** statement, see the **if** command.

The **passes-through** operator takes a sequence of integers or parameters, enclosed in single quotation marks, as an argument. It can also take a single integer or parameter as an argument. It evaluates to true if the supplied integer or parameter appears anywhere in the AS path, or if the supplied sequence of integers and parameters appears, in the same order, anywhere in the AS path. This includes the **originates-from** or **neighbor-is** location in the AS path.

Task ID

Task ID	Operations
route-policy	read, write

Examples

The following are incomplete configuration examples:

```
RP/0/RSP0/CPU0:router(config-rpl)# if as-path passes-through '10' then
RP/0/RSP0/CPU0:router(config-rpl-if)# if as-path passes-through '$asnum' then
RP/0/RSP0/CPU0:router(config-rpl-if)# if as-path passes-through '10 11' then
RP/0/RSP0/CPU0:router(config-rpl-if)# if as-path passes-through '10 $asnum 12' then
```

Without the **exact** keyword, repeated autonomous system numbers in the AS path are ignored. For example:

```
RP/0/RSP0/CPU0:router(config-rpl)# if as-path passes-through '9 10 11' then
```

matches an AS path containing

```
...9 10 10 10 11 ....
```

and an AS path containing:

```
...9 10 11...
```

With the **exact** keyword, repetitions are not ignored. Therefore:

```
RP/0/RSP0/CPU0:router(config-rpl)# if as-path passes-through '9 10 11' exact then
```

matches the second of these AS paths but not the first.

Related Commands

Command	Description
as-path in, on page 1605	Matches the AS path of a route to an AS path set.
as-path length, on page 1609	Compares the number of ASN in the AS path of a route
as-path originates-from, on page 1614	Compares an AS path to the AS sequence beginning with the AS number that originated a route.
as-path unique-length, on page 1622	Performs specific checks based on the length of the AS path.

as-set

To create a named AS set, use the **as-set** command in global configuration mode. To remove the named AS set, use the **no** form of this command. **as-set** accelerates the look-up time.

Note **as-set** is faster than **as-path-set** if the only requirement is to check the origin of AS.

as-set *name*

Syntax Description

name	Name of the AS path set.

Command Default

No default behavior or values

Command Modes

global configuration

Command History

Release	Modification
Release 7.3.1	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. Use the **as-set** command to create a named AS set. This command enters AS set configuration mode, in which you can use any of the below option to specify an operation. This as-set can be used only with "as-path originates-from", please refer example below.

Options	Description
<1-4294967295>	32 bit number indicating the autonomous system (AS).
<1-65535>	16 bit number indicating the autonomous system (AS).

Task ID

Task ID	Operations
route-policy	read, write

The following is a sample definition of an AS set named set1. This AS set is composed of two elements. This **as-set** can be used only with **as-path originates-from** in the route-policy as shown in the example below:

```
as-set set1
 1240,
1.2,
1568,
2345,
 2.3,
```

```
 1297,
end-set

route-policy sample
if as-path originates-from set1 then
  pass
endif
end-policy

Incoming Route 1: AS-path= 123,457,895,443, 1568   For this route the condition evaluates
to 'TRUE' because 1568 is there in as-set "set1"
Incoming Route 2: AS-path= 123,457,895,596  AS 596 is not there in as-set "set1" , the
condition evaluates to FALSE
Incoming  Route 3: AS-path= 123,457,895,443,100  As 100 is not there in as-set "set1" , the
 condition evaluates to FALSE
Incoming Route 4: AS-path= 123,457,895,443,1240  For this route the condition evaluates to
 'TRUE' because 1240 is there in as-set "set1"
```

as-path-set

To create a named AS path set, use the **as-path-set** command in global configuration mode. To remove the named AS path set, use the **no** form of this command.

Note **as-set** faster than **as-path-set** if the only requirement is to check the origin of AS. For more information, see as-set, on page 1618.

as-path-set *name*
no as-path-set *name*

Syntax Description

name	Name of the AS path set.

Command Default

No default behavior or values

Command Modes

global 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **as-path-set** command to create a named AS path set.

An AS path set comprises operations for matching an AS path attribute.

This command enters AS path set configuration mode, in which you can use any of the below option to specify an operation.

Options	Description
dfa-regex	Indicates the DFA (deterministic finite automata) style regular expression. It performs better for complex regular expressions. Single quotation marks are required around the regular expression.
ios-regex	Indicates the traditional IOS style regular expression. It performs better with simpler regular expressions. Single quotation marks are required around the regular expression.
length	Indicates the number of ASN (Autonomous System Number) in the AS path of a Border Gateway Protocol (BGP) route.

Options	Description
neighbor-is	Indicates the neighbor's AS-path number that can be matched with.
originates-from	Indicates the BGP AS from which the route originated.
passes-through	Indicates if the supplied integer or parameter appears anywhere in the AS path, or if the supplied sequence of integers and parameters appear, in the same order, anywhere in the AS path.
unique-length	Indicates the length of BGP AS-path, ignoring duplicates.

The above options can also be used as an inline set in a parenthesized list of comma-separated expressions.

Task ID

Task ID	Operations
route-policy	read, write

Examples

The following is a sample definition of an AS path set named aset1. This AS path set is composed of two elements. When used in a matching operation, this AS path set matches any route whose AS path ends with either the autonomous system number 42 or 127.

```
RP/0/RSP0/CPU0:router(config)# as-path-set aset1
RP/0/RSP0/CPU0:router(config-as)# ios-regex '_42$',
RP/0/RSP0/CPU0:router(config-as)# ios-regex '_127$'
RP/0/RSP0/CPU0:router(config-as)# end-set
```

The following is a sample of the as-path options used as an inline set.

```
RP/0/RSP0/CPU0:router(config-rpl)# if as-path in (ios-regex '_42$', ios-regex$ '_127$')
RP/0/RSP0/CPU0:router(config-rpl-if)# pass
RP/0/RSP0/CPU0:router(config-rpl-if)# endif
RP/0/RSP0/CPU0:router(config-rpl)#
```

as-path unique-length

To perform specific checks based on the length of the AS path (match against the number of unique ASNs in the AS path), use the **as-path unique-length** command in route-policy configuration mode.

as-path unique-length {**eq** | **is** | **ge** | **le**} {*numberparameter*}

Syntax Description

eq | **is** | **ge** | **le**	Equal to; greater than or equal to; less than or equal to.
number	Value assigned to an 11-bit unsigned integer. Range is from 0 to 2047.
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 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **as-path unique-length** command as a conditional expression within an **if** statement to perform a match based on the length of the AS path.

Note For a list of all conditional expressions available within an **if** statement, see the **if** command.

The **unique-length** operator is similar to the length operator, except that when an AS path has been padded with the same autonomous system number multiple times, the operator counts only one when the route is padded. Therefore, given an AS path of 333 333 111 222 123 444 444 444, the **unique-length** operator would return a value of 5, whereas the length operator would return a value of 8.

Task ID

Task ID	Operations
route-policy	read, write

Examples

The following example shows show to perform checks based on the AS path length. If the AS path matches the specified values, the local preference is set to 100:

```
RP/0/RSP0/CPU0:router(config-rpl)# if as-path unique-length eq 10 then
RP/0/RSP0/CPU0:router(config-rpl-if)# if as-path unique-length ge 10 then
RP/0/RSP0/CPU0:router(config-rpl-if)# if as-path unique-length le 10 then
```

```
RP/0/RSP0/CPU0:router(config-rpl)# if as-path unique-length eq $integerparam then
RP/0/RSP0/CPU0:router(config-rpl-if)# if as-path unique-length ge $geparam then
RP/0/RSP0/CPU0:router(config-rpl-if)# if as-path unique-length le $leparam then

RP/0/RSP0/CPU0:router(config-rpl)# set local-preference 100
RP/0/RSP0/CPU0:router(config-rpl)# endif
```

Related Commands

Command	Description
as-path length, on page 1609	Performs conditional checks based on the length of the AS path.

community is-empty

To check if a route has no community attributes associated with it, use the **community is-empty** command in route-policy configuration mode.

community is-empty

Syntax Description

This command has no arguments or keywords.

Command Default

No default behavior or values

Command Modes

Route-policy 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **community is-empty** command as a conditional expression within an **if** statement to check if a route has community attributes associated with it.

Note For a list of all conditional expressions available within an **if** statement, see the **if** command.

This command takes no arguments and evaluates to true only if the route has no community attributes associated with it.

Task ID

Task ID	Operations
route-policy	read, write

Examples

In the following example, if the route has no community attributes associated with it, then the local preference is set to 100:

```
RP/0/RSP0/CPU0:router(config-rpl)# if community is-empty then
RP/0/RSP0/CPU0:router(config-rpl-if)# set local-preference 100
RP/0/RSP0/CPU0:router(config-rpl-if)# endif
```

community matches-any

To match any elements of a community set, use the **community matches-any** command in route-policy configuration mode.

community matches-any {*community-set-nameinline-community-setparameter*}

Syntax Description

community-set-name	Name of a community set.
inline-community-set	Inline community set. The inline community set must be enclosed in parentheses.
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 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **community matches-any** command as a conditional expression within an **if** statement to match any element of a community set.

Note For a list of all conditional expressions available within an **if** statement, see the **if** command.

A simple condition using the **matches-any** operator evaluates as true if at least one community element of the community attribute for the route matches an element in the community set operand. If no community in the route matches any of the specifications in the named or inline set, then the condition evaluates to false. Likewise, when there is no community at all in the route, the condition evaluates to false.

Matching of a community in the route to a specification in a named or an inline set is intuitive. If the community specification in a set is the familiar colon-separated decimal 16-bit numbers specification, or one of the well-known communities, the community matches the specification if the specification denotes the same 32-bit number as that in the route. If the community specification uses a wildcard, then the community in the route matches if it is one of the many communities denoted by the wildcard specification. In inline sets, community specifications may be parameterized, in which case the relevant matching is done when the value of the parameter has been supplied.

Communities may also be matched using range and regular expression operators. Range specifications are entered as follows: [*low-value .. high-value*]. Either or both colon-separated halves of a community value may contain a range. The following are valid range specifications:

```
10:[100..1000]
```

```
[10..100]:80
[10..100]:[100..2000]
```

In addition, the **private-as** keyword may be used to specify the range from 64512 to 65534. Regular expressions are specified as the **ios-regex** keyword followed by a valid regular expression string.

Community values from the route are matched one at a time to the match specifications. Therefore, regex match specifications are expected to represent one individual community value and not a sequence of community values.

Task ID

Task ID	Operations
route-policy	read, write

Examples

In the following example, a named community set called my-community-set and a route policy called community-matches-any-example are created. The policy sets the local-preference to 100 for any route that has one or more of the communities in the my-community-set community set. If the route does not have any of these communities, the policy checks whether it has any communities whose first half is in the range from 10 to 25 and whose second half is the value 35, in which case it sets the local-preference to 200. Otherwise, it checks for a community value in the range of 30:100 to 30:500, in which case it sets the local-preference to 300.

```
RP/0/RSP0/CPU0:router(config)# community-set my-community-set
RP/0/RSP0/CPU0:router(config-comm)# 10:20,
RP/0/RSP0/CPU0:router(config-comm)# 10:30,
RP/0/RSP0/CPU0:router(config-comm)# 10:40
RP/0/RSP0/CPU0:router(config-comm)# end-set

RP/0/RSP0/CPU0:router(config)# route-policy community-matches-any-example
RP/0/RSP0/CPU0:router(config-rpl)# if community matches-any my-community-set then
RP/0/RSP0/CPU0:router(config-rpl-if)# set local-preference 100
RP/0/RSP0/CPU0:router(config-rpl-if)# elseif community matches-any ([10..25]:35) then
RP/0/RSP0/CPU0:router(config-rpl-elseif)# set local-preference 200
RP/0/RSP0/CPU0:router(config-rpl-elseif)# elseif community matches-any (30:[100..500])
then
RP/0/RSP0/CPU0:router(config-rpl-elseif)# set local-preference 300
RP/0/RSP0/CPU0:router(config-rpl-elseif)# endif
RP/0/RSP0/CPU0:router(config-rpl)# end-policy
```

Related Commands

Command	Description
community matches-every, on page 1627	Matches every element of a community set.

community matches-every

To match every element of a community set, use the **community matches-every** command in route-policy configuration mode.

community matches-every {*community-set-nameinline-community-setparameter*}

Syntax Description

community-set-name	Name of a community set.
inline-community-set	Inline community set. The inline community set must be enclosed in parentheses.
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 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **community matches-every** command as a conditional expression within an **if** statement to match every element of a community set.



Note For a list of all conditional expressions available within an **if** statement, see the **if** command.

A simple condition using the **matches-every** operator evaluates as true if every specification in the named set or inline set specified matches at least one community value in the route. If any community specification in the named or inline set is not matched, then the operation evaluates to false.

Matching of a community in the route to a specification in a named or an inline set is intuitive. If the community-specification in a set is the familiar colon-separated decimal 16-bit numbers specification, or one of the well-known communities, the community matches the specification if the specification denotes the same 32-bit number as that in the route. If the community specification uses a wildcard, then the community in the route matches if it is one of the many communities denoted by the wildcard specification. In inline sets, community specifications may be parameterized, in which case the relevant matching is done when the value of the parameter has been supplied.

Communities may also be matched using range and regular expression operators. Range specifications are entered as follows: [*low-value .. high-value*]. Either or both colon-separated halves of a community value may contain a range. The following are valid range specifications:

```
10:[100..1000]
[10..100]:80
```

```
[10..100]:[100..2000]
```

Therefore, a **matches-every** operation with two community range specifications means that a community must be present in the route that corresponds to each range. For example, in the following statement:

```
if community matches-every (10:[100..200],20:[100..200]) then
```

the statement evaluates as true if one or more communities in the route lie in the range 10:[100.200] and one or more communities in the route lie in the range 20:[100..200].

In addition, the **private-as** keyword may be used to specify the range from 64512 to 65534.

Regular expressions are specified as the **ios-regex** keyword followed by a valid single-quoted regular expression string. Community values from the route are matched one at a time against the match specifications. Therefore, regex match specifications are expected to represent one individual community value and not a sequence of community values.

Task ID

Task ID	Operations
route-policy	read, write

Examples

In the following example, the route policy named community-matches-every-example sets the local-preference value to 100 for all routes that have all three communities in the my-community-set community set. Routes that do not have all three communities but have a community that matches the first regular expression match have the local-preference value set to 200. Finally, any remaining routes that match the last regular expression have the local-preference values set to 300.

```
  RP/0/RSP0/CPU0:router(config)# community-set my-community-set
  RP/0/RSP0/CPU0:router(config-comm)# 10:20,
  RP/0/RSP0/CPU0:router(config-comm)# 10:30,
  RP/0/RSP0/CPU0:router(config-comm)# 10:40
  RP/0/RSP0/CPU0:router(config-comm)# end-set

  RP/0/RSP0/CPU0:router(config)# route-policy community-matches-every-example
  RP/0/RSP0/CPU0:router(config-rpl)# if community matches-every my-community-set then
  RP/0/RSP0/CPU0:router(config-rpl-if)# set local-preference 100
  RP/0/RSP0/CPU0:router(config-rp-elseif)# elseif community matches-every (ios-regex
'_10:[0-9]0_') then
  RP/0/RSP0/CPU0:router(config-rpl-elseif)# set local-preference 200
  RP/0/RSP0/CPU0:router(config-rpl-elseif)# elseif community matches-every
(ios-regex'_20:[0-9]0_') then
  RP/0/RSP0/CPU0:router(config-rpl-elseif)# set local-preference 300
  RP/0/RSP0/CPU0:router(config-rpl-elseif)# endif
  RP/0/RSP0/CPU0:router(config-rpl)# end-policy
```

Related Commands

Command	Description
community matches-any, on page 1625	Matches any element of a community set.

community matches-within

To configure a route policy to match within a community set, use the **community matches-within** command in route-policy configuration mode.

```
community matches-within { community-set-name or inline-community-set | parameter }
```

Syntax Description

community-set-name	Name of a community set.
inline-community-set	Inline community set. The inline community set must be enclosed in parentheses.
parameter	Parameter name. The parameter name must be preceded with a "$."

Command Default

No default behavior or values

Command Modes

Route-policy configuration

Task ID

Task ID	Operations
route-policy	read, write

Command History

Release	Modification
Release 6.3.1	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.

This command is similar to the **community matches-any** command, but every community in the route must match at least one match specification. If the route has no communities, then it matches.

Note For a list of all conditional expressions available within an **if** statement, see the **if** command.

Task ID

Task ID	Operations
route-policy	read, write

The following example shows how to configure a route policy to match within the elements of a community set.

```
RP/0/RSP0/CPU0:router#config
RP/0/RSP0/CPU0:router(config)#route-policy bob
RP/0/RSP0/CPU0:router(config-rpl)#if community matches-within (*:3, 5:*) then
```

```
RP/0/RSP0/CPU0:router(config-rpl)#set local-preference 94
RP/0/RSP0/CPU0:router(config-rpl)#endif
RP/0/RSP0/CPU0:router(config-rpl)#end-policy
```

For example, routes with these sets of communities return TRUE:

- (1:3, 5:10)
- (5:3)
- (2:3, 6:3, 4:3)

Routes with the following set of communities return FALSE:

(1:3, 5:10, 6:5) —The community (6:5) does not match

community-set

To define a community set, use the **community-set** command in global configuration mode. To remove the community set, use the **no** form of this command.

community-set *name*
no community-set *name*

Syntax Description

name	Name of the community set.

Command Default

No default behavior or values

Command Modes

global 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Regular expressions and ranges can be specified to match the communities. An attempt to use a community set that contains a range or regular expression to set a community value is rejected when an attempt to attach such a policy is made.

A community set holds community values for matching against the Border Gateway Protocol (BGP) community attribute. A community is a 32-bit quantity. For notational convenience, each community value must be split in half and expressed as two unsigned decimal integers in the range from 0 to 65535, separated by a colon.

The inline form of a community set also supports parameterization. Each 16-bit portion of the community may be parameterized.

The routing policy language (RPL) provides symbolic names for the standard well-known community values: **accept-own** is 0xFFFF0001, **internet** is 0:0, **no-export** is 65535:65281, **no-advertise** is 65535:65282, and **local-as** is 65535:65283.

RPL also provides a facility for using wildcards in community specifications. A wildcard is specified by inserting an asterisk (*) in place of one of the 16-bit portions of the community specification, which indicates that any value for that portion of the community matches.

Every community set must contain at least one community value. An empty community set is invalid and the policy configuration system rejects it.

Community sets can be entered in these formats:

Format	Description
#-remark	Remark beginning with '#'
*	Wildcard (any community or part thereof)
0-65535	16-bit half-community number

Format	Description
[	Left bracket to begin range
accept-own	Accept-Own (BGP well-known community)
dfa-regex	DFA (deterministic finite automata) style regular expression
internet	Internet (BGP well-known community)
ios-regex	Traditional IOS style regular expression
local-AS	Do not send outside local AS (BGP well-known community)
no-advertise	Do not advertise to any peer (BGP well-known community)
no-export	Do not export to next AS (BGP well-known community)
private-as	Match within BGP private AS range [64512..65534]



Note The dfa-regex and ios-regex syntax for community set is *"[']*[^':&<>]*:*[^':&<>]*[']"*. This means that regex starts with a single-quote (') followed by a string of any character (that does not include single-quote, colon, ampersand, less-than, greater-than, or space) followed by a colon, and a string of any characters (that does not include single-quote, colon, ampersand, less-than, greater-than, or space) followed by single-quote.

Task ID

Task ID	Operations
route-policy	read, write

Examples

In the following example, a community set named cset_accept_own is created:

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#community-set cset_accept_own
RP/0/RSP0/CPU0:router(config-comm)#accept-own
RP/0/RSP0/CPU0:router(config-comm)#end-set
```

In the following example, a community set named cset1 is created:

```
RP/0/RSP0/CPU0:router(config)# community-set cset1
RP/0/RSP0/CPU0:router(config-comm)# 12:34,
RP/0/RSP0/CPU0:router(config-comm)# 12:56,
RP/0/RSP0/CPU0:router(config-comm)# 12:78,
RP/0/RSP0/CPU0:router(config-comm)# internet
RP/0/RSP0/CPU0:router(config-comm)# end-set
```

In the following example, a community set named cset2 is created:

```
RP/0/RSP0/CPU0:router(config)# community-set cset2
RP/0/RSP0/CPU0:router(config-comm)# 123:456,
RP/0/RSP0/CPU0:router(config-comm)# no-advertise,
RP/0/RSP0/CPU0:router(config-comm)# end-set
```

In the following example, a community set named cset3 is created. This policy uses wildcards and matches all communities where the autonomous system part of the community is 123.

```
RP/0/RSP0/CPU0:router(config)# community-set cset3
RP/0/RSP0/CPU0:router(config-comm)# 123:*
RP/0/RSP0/CPU0:router(config-comm)# end-set
```

delete community

To delete community attributes associated with a Border Gateway Protocol (BGP) route, use the **delete community** command in route-policy configuration mode.

delete community {**all** | **in** {*community-set-nameinline-community-setparameter*} | **not in** {*community-set-nameinline-community-setparameter*}}

Syntax Description

all	Removes all communities except the well-known communities.
in	Removes any communities associated with the route that are listed in either the named community set or the inline community set.
community-set-name	Name of a community set.
inline-community-set	Inline community set. The inline community set must be enclosed in parentheses.
parameter	Parameter name. The parameter name must be preceded with a "$."
not in	Removes all communities that are not listed in either the named community set or the inline community set, and are not well-known communities.

Command Default

No default behavior or values

Command Modes

Route-policy 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **delete community** command to delete community attributes associated with a BGP route.

Note The **delete community** command can be used as an action statement within an **if** statement. For a list of all action statements available within an **if** statement, see the **if** command.

Communities are 32-bit values carried in Border Gateway Protocol (BGP) routes. Each route may have zero or more communities in an unordered list.

You can remove a well-known community (internet, no-export, no-advertise, or local-as) from a route, but this removal must be done explicitly. This command should be used with a degree of caution. In general, few circumstances exist in which you would need to remove a well-known community.

Task ID

Task ID	Operations
route-policy	read, write

Examples

The following example shows how to delete any communities associated with the routes that are listed in either the named community set or inline community set, respectively.

```
RP/0/RSP0/CPU0:router(config-rpl)# delete community in my_community_set
RP/0/RSP0/CPU0:router(config-rpl)# delete community in (10:[0..50],20:[60..80])
```

The following example shows how to remove all communities including well-known communities.

```
RP/0/RSP0/CPU0:router(config-rpl)# delete community in (internet, no-export, no-advertise,
local- as, *:*)
```

The following example shows how to remove all communities except for the well-known communities.

```
RP/0/RSP0/CPU0:router(config-rpl)# delete community all
```

The following example shows how to delete the well-known community value internet from a route:

```
RP/0/RSP0/CPU0:router(config-rpl)# delete community in (internet)
```

delete large-community

To delete the specified large-communities from a route policy, use the **delete large-community** command in the route-policy configuration mode.

```
delete large-community { all |[not] in { named or inline-large-community-set |
parameter } }
```

Syntax Description

all	Removes all large communities.
in	Removes any large communities associated with the route that are listed in either the named large community set or the inline large community set.
large-community-set-name	Name of a large community set.
inline-large-community-set	Inline large community set. The inline community set must be enclosed in parentheses.
parameter	Parameter name. The parameter name must be preceded with a "$."
not	Removes all communities that are not listed in either the named large community set or the inline large community set.

Command Default

No default behavior or values

Command Modes

Route-policy configuration

Command History

Release	Modification
Release 6.3.1	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.

The large communities 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.

In route-policy statements, each integer in the BGP large community can be replaced by any of the following expressions:

- [x..y] — This expression specifies a range between x and y, inclusive.
- * — This expression stands for any number.
- peeras — This expression is replaced by the AS number of the neigbhor from which the community is received or to which the community is sent, as appropriate.
- not-peeras — This expression matches any number other than the peeras.
- private-as — This expression specifies any number in the private ASN range: [64512..65534] and [4200000000..4294967294].

Note The peeras and not-peeras expressions can only be used in delete statements that appear in route policies that are applied at the neighbor-in or neighbor-out attach points.

IOS regular expression (ios-regex) and DFA style regular expression (dfa-regex) can be used in the delete statements. For example, the IOS regular expression ios-regex '^5:.*:7$' is equivalent to the expression 5:*:7.

Task ID

Task ID	Operations
route-policy	read, write

The following example shows how to delete specified BGP large-communities from a route policy using the delete large-community command.

```
RP/0/RP/0/RSP0/CPU0:router#config
RP/0/RSP0/CPU0:router(config)#route-policy lrg_comm_rp2S
RP/0/RSP0/CPU0:router(config-rpl)#delete large-community in (ios-regex '^100000:')
RP/0/RSP0/CPU0:router(config-rpl)#delete large-community not in (peeras:*:*, 41289:*:*)
RP/0/RSP0/CPU0:router(config-rpl)#delete large-community in catbert
RP/0/RSP0/CPU0:router(config-rpl)#end-policy
```

delete extcommunity rt

To delete route target (RT) extended community attributes associated with a Border Gateway Protocol (route), use the **delete extcommunity rt** command in route-policy configuration mode.

delete extcommunity rt {**all** | **in** {*extcommunity-set-nameinline-extcommunity-setparameter*} | **not in** {*extcommunity-set-nameinline-extcommunity-setparameter*}}

Syntax Description

all	Removes all extended communities.
in	Removes any extended communities associated with the routes that are listed in either the named extended community set or the inline extended community set.
extcommunity-set-name	Name of an extended community set.
inline-extcommunity-set	Inline extended community set. The inline extended community set must be enclosed in parentheses.
parameter	Parameter name. The parameter name must be preceded with a "$."
not in	Removes all extended communities that are not listed in either the named extended community set or the inline extended community set, and are not well-known extended communities.

Command Default

No default behavior or values

Command Modes

Route-policy 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **delete extcommunity rt** command to delete extended community values from a BGP route target extended community list in a route.

Note The **delete extcommunity rt** command can be used as an action statement within an **if** statement. For a list of all action statements available within an **if** statement, see the **if** command.

Extended communities are similar to regular Border Gateway Protocol (BGP) communities but contain more data and have a richer structure for encoding information in them.

Extended communities can be in the following forms: SoO:AS:tag, SoO:IP:tag, RT:AS:tag, or RT:IP:tag.

Wildcards (*) and regular expressions are allowed for extended community set elements.

The forms of this command that take a named extended community set or an inline extended community set value as arguments are equivalent. They delete any extended communities that are listed in either the named set or the inline set, respectively.

Task ID

Task ID	Operations
route-policy	read, write

Examples

In the following example, all extended communities are deleted:

```
RP/0/RSP0/CPU0:router(config-rpl)# delete extcommunity rt all
```

In this example, any extended communities that are listed in my-extcommunity-set are deleted:

```
RP/0/RSP0/CPU0:router(config-rpl)# delete extcommunity rt in my-extcommunity-set
```

In this example, extended communities associated with the route listed in the named inline extended community sets are deleted:

```
RP/0/RSP0/CPU0:router(config-rpl)# delete extcommunity rt in (67:29, 67:55)
```

destination in

To match a destination entry in a named prefix set or inline prefix set, use the **destination in** command in route-policy configuration mode.

destination in {*prefix-set-nameinline-prefix-setparameter*}

Syntax Description

prefix-set-name	Name of a prefix set.
inline-prefix-set	Inline prefix set. The inline prefix set must be enclosed in parentheses.
parameter 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 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Note

- Use the **destination in** command as a conditional expression within an **if** statement to match a destination entry in a named prefix set or inline prefix set.
- For a list of all conditional expressions available within an **if** statement, see the **if** command.
- This command takes either a named prefix set or an inline prefix set value as an argument. The condition returns true if the destination entry matches any entry in the prefix set or inline prefix set. An attempt to match a destination using a prefix set that is defined but contains no elements returns false.
- The routing policy language (RPL) provides the ability to test destinations for a match to a list of prefix match specifications using the **in** operator. The **destination in** command is protocol-independent.
- Use the **destination-prefix in** command to filter a flowspec route in BGP, and the **match destination-address** command to create a flowspec rule that matches a specific destination.
- Use the **destination in** command to filter BGP routes for other address families, and use the **destination-prefix in** command to filter flowspec routes in BGP that include a **match destination-address** clause.
- RPL supports both 32-bit IPv4 prefixes, specified in dotted-decimal format, and 128-bit IPv6 prefixes, specified in colon-separated hexadecimal format.

Task ID

Task ID	Operations
route-policy	read, write

Examples

In the following example, a prefix set named my-prefix-set is defined and a route policy named use-destination-in is created. Within the use-destination-in route policy, the **destination in** command is used within an **if** statement to learn if the destination is in the prefix-set named my-prefix-set. If it is, then local preference is set to 100. If it is not in my-prefix-set but does match the next prefix specifications, then local preference is set to 200.

```
RP/0/RSP0/CPU0:router(config)# prefix-set my-prefix-set
RP/0/RSP0/CPU0:router(config-pfx)# 10.0.0.1/32,
RP/0/RSP0/CPU0:router(config-pfx)# fe80::203:0:0:0/64,
RP/0/RSP0/CPU0:router(config-pfx)# 10.0.0.2/24 le 32
RP/0/RSP0/CPU0:router(config-pfx)# end-set

RP/0/RSP0/CPU0:router(config)# route-policy use-destination-in
RP/0/RSP0/CPU0:router(config-rpl)# if destination in my-prefix-set then
RP/0/RSP0/CPU0:router(config-rpl-if)# set local-preference 100
RP/0/RSP0/CPU0:router(config-rpl-if)# elseif destination in (10.0.0.1/32, 10.0.0.2/24 le
32) then
RP/0/RSP0/CPU0:router(config-rpl-elseif)# set local-preference 200
RP/0/RSP0/CPU0:router(config-rpl-elseif)# endif
RP/0/RSP0/CPU0:router(config-rpl)# end-policy
```

In the following example, a prefix set named ipv6-prefix-set is defined and a route policy named ipv6-destination-in is created. Within the ipv6-destination-in route policy, the **destination in** command is used within an **if** statement to learn if the destination is in the prefix-set named ipv6-prefix-set. If it is, then the next-hop is set to 2001:abcd:fedc::1. If it is not in ipv6-prefix-set but does match the next prefix specifications, then the next-hop is set to 1111:2222:3333:4444:5555:6666:7777:8888.

```
RP/0/RSP0/CPU0:router(config)# prefix-set ipv6-prefix-set
RP/0/RSP0/CPU0:router(config-pfx)# 2001:0:0:1::/64,
RP/0/RSP0/CPU0:router(config-pfx)# 2001:0:0:2::/64,
RP/0/RSP0/CPU0:router(config-pfx)# 2001:0:0:3::/64,
RP/0/RSP0/CPU0:router(config-pfx)# 2001:0:0:4::/64
RP/0/RSP0/CPU0:router(config-pfx)# end-set

RP/0/RSP0/CPU0:router(config)# route-policy ipv6-destination-in
RP/0/RSP0/CPU0:router(config-rpl)# if destination in ipv6-prefix-set then
RP/0/RSP0/CPU0:router(config-rpl-if)# set next-hop 2001:abcd:fedc::1
RP/0/RSP0/CPU0:router(config-rpl-if)# elseif destination in (2001::1, 2002:1:2:3::/64)
then
RP/0/RSP0/CPU0:router(config-rpl-elseif)# set next-hop
1111:2222:3333:4444:5555:6666:7777:8888
RP/0/RSP0/CPU0:router(config-rpl-elseif)# endif
RP/0/RSP0/CPU0:router(config-rpl)# end-policy
```

done

To stop executing a policy and accept the route, use the **done** command in route-policy configuration mode.

done

Syntax Description

This command has no arguments or keywords.

Command Default

No default behavior or values

Command Modes

Route-policy 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **done** command to stop executing the policy and accept the route.

Note The **done** command can be used as an action statement within an **if** statement. For a list of all action statements available within an **if** statement, see the **if** command.

When encountering a **done** statement the route is passed and no further policy statements are executed. All modifications made to the route prior to the **done** statement are still valid.

Note The default action of a route policy is to drop or discard any routes that have not been either explicitly passed or for which no attempt has been made to modify with an action. The routing policy language (RPL) does not have specific "match clauses," which means the default drop behavior is controlled by whether a route has been explicitly passed or an attempt has been to modify the route using an action statement.

Task ID

Task ID	Operations
route-policy	read, write

Examples

In the following example, if the destination match succeeds for 29.0.0.0/8 le 32, the execution continues past set community 102:12 and onto the next statement. If the destination match succeeds for 39.0.0.0/8 le 32 execution, then the policy execution stops when in encounters the *done* statement.

```
RP/0/RSP0/CPU0:router(config)# route-policy done_st_example
RP/0/RSP0/CPU0:router(config-rpl)# if destination in (29.0.0.0/8 le 32) then
```

```
RP/0/RSP0/CPU0:router(config-rpl-if)# set community 102:12
RP/0/RSP0/CPU0:router(config-rpl-if)# endif
RP/0/RSP0/CPU0:router(config-rpl)# if destination in (39.0.0.0/8 le 32) then
RP/0/RSP0/CPU0:router(config-rpl-if)# set community 102:39
RP/0/RSP0/CPU0:router(config-rpl-if)# done
RP/0/RSP0/CPU0:router(config-rpl-if)# endif
RP/0/RSP0/CPU0:router(config-rpl)# if destination in (49.0.0.0/8 le 32) then
RP/0/RSP0/CPU0:router(config-rpl-if)# set community 102:49
RP/0/RSP0/CPU0:router(config-rpl-if)# endif
RP/0/RSP0/CPU0:router(config-rpl)# if destination in (59.0.0.0/8 le 32) then
RP/0/RSP0/CPU0:router(config-rpl-if)# set community 102:59
RP/0/RSP0/CPU0:router(config-rpl-if)# endif
RP/0/RSP0/CPU0:router(config-rpl)# end-policy
```

drop

To discard a route, use the **drop** command in route-policy configuration mode.

drop

Syntax Description

This command has no arguments or keywords.

Command Default

No default behavior or values

Command Modes

Route-policy 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **drop** command within a route policy to drop a route.

Note The **drop** command can be used as an action statement within an **if** statement. For a list of all action statements available within an **if** statement, see the **if** command.

This command causes the route to be dropped. After a route is dropped, no further execution of policy occurs. Therefore, if after executing the first two statements of a policy the **drop** statement is encountered, the route is discarded and execution stops immediately even when the policy contains further statements.

Note The default action of a route policy is to drop or discard any routes that have not been either explicitly passed or attempted to be modified with an action. The routing policy language (RPL) does not have specific "match clauses," which means the default drop behavior is controlled by whether a route has been explicitly passed or an attempt has been to modify the route using an action statement.

Task ID

Task ID	Operations
route-policy	read, write

Examples

In the following example, any route with a destination address contained within the prefix set pset1 is dropped:

```
RP/0/RSP0/CPU0:router(config-rpl)# if destination in pset1 then
RP/0/RSP0/CPU0:router(config-rpl-if)# drop
```

```
RP/0/RSP0/CPU0:router(config-rpl-if)# endif
RP/0/RSP0/CPU0:router(config-rpl)#
```

edit

To edit the contents of a route policy, a prefix set, an AS path set, a community set, or an extended community set, use the **edit** command in EXEC configuration mode.

edit {**route-policy** | **prefix-set** | **as-path-set** | **community-set** | **extcommunity-set** {**rt** | **soo**} | **policy-global** | **rd-set**} *name* [**nano** | **emacs** | **vim** | **inline** {**add** | **prepend** | **remove**} *set-element*]

Syntax Description

route-policy	Edits the contents of a route policy.
prefix-set	Edits the contents of a prefix set.
as-path-set	Edits the contents of an AS path set.
community-set	Edits the contents of a community set.
extcommunity-set	Edits the contents of an extended community set of the specified type.
rt	Edits the BGP route target (RT) extended community.
soo	Edits the BGP site of origin (SoS) extended community.
policy-global	Edits the contents of policy-global definitions.
rd-set	Edits the contents of a route-distinguisher set.
name	Name of a route policy, a prefix set, an AS path set, a community set, or an extended community set, RD set, or global parameters.
nano	(Optional) Uses GNU Nano text editor.
emacs	(Optional) Uses Micro Emacs editor.
vim	(Optional) Uses VI Improved editor.
inline	(Optional) Uses the command line.
add	Appends the element to the set.
prepend	Prepends the element to the set.
remove	Removes the element from the set.
set-element	Value of the set element. **Note** To inline edit multiple set elements separated with comma, use quotes to club the entries as a single argument. Example: `edit extcommunity-set rt rt_set inline add "4:4,5:4"`

Command Default

Default editor is GNU vim text editor

Command Modes

EXEC configuration

Command History

Release	Modification
Release 3.7.2	This command was introduced.
Release 7.11.1	The **Nano** and **Emacs** keyword was deprecated.

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.

Use the **edit** command to edit the contents of a route policy, a prefix set, an AS path set, a community set, an extended community set, a global policy, or a route destination set.

After editing with Nano, save the edit buffer and exit the editor using the Ctrl-X keystroke.

After editing with Emacs, save the editor buffer by using the Ctrl-X and Ctrl-S keystrokes. To save and exit the editor, use the Ctrl-X and Ctrl-C keystrokes.

After editing with VIM, to write to a current file and exit use the :wq or :x or ZZ keystrokes. To quit and confirm, use the :q keystrokes. To quit and discard changes, use the :q! keystrokes.

Task ID

Task ID	Operations
route-policy	read, write

Examples

In the following example, the policy_A policy is opened in the editor:

```
RP/0/RSP0/CPU0:router# edit route-policy policy_A

----------------------------------------
== MicroEMACS 3.8b () == rpl_edit.139281 ==
  if destination in (2001::/8) then
    drop
  endif
end-policy
!

== MicroEMACS 3.8b () == rpl_edit.139281 ==
Parsing.
83 bytes parsed in 1 sec (82)bytes/sec
Committing.
1 items committed in 1 sec (0)items/sec
Updating.
Updated Commit database in 1 sec
```

If there are parse errors, you are asked whether editing should continue:

```
RP/0/RSP0/CPU0:router#edit route-policy policy_B
== MicroEMACS 3.8b () == rpl_edit.141738
route-policy policy_B
 set metric-type type_1
```

```
 if destination in (2001::/8) then
    drop
  endif
end-policy
!
== MicroEMACS 3.8b () == rpl_edit.141738 ==
Parsing.
105 bytes parsed in 1 sec (103)bytes/sec

% Syntax/Authorization errors in one or more commands.!! CONFIGURATION
FAILED DUE TO SYNTAX/AUTHORIZATION ERRORS
 set metric-type type_1
 if destination in (2001::/8) then
    drop
  endif
end-policy
!

Continue editing? [no]:
```

If you answer **yes** , the editor continues on the text buffer from where you left off. If you answer **no**, the running configuration is not changed and the editing session is ended.

After the policy is opened, it may be manipulated using normal editor commands, then saved and committed to the running configuration.

end-global

To end the definition of global parameters and exit global parameter configuration mode, use the **end-global** command in global parameter configuration mode.

end-global

Syntax Description

This command has no arguments or keywords.

Command Default

No default behavior or values

Command Modes

Global parameter 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **end-global** command to end the definition of global parameters and exit global parameter configuration mode.

Task ID

Task ID	Operations
route-policy	read, write

Examples

In the following example, the **end-global** command ends the definition of global parameters:

```
RP/0/RSP0/CPU0:router(config)#policy-global
RP/0/RSP0/CPU0:router(config-rp-gl)# glbpathtype 'ebgp'
RP/0/RSP0/CPU0:router(config-rp-gl)# glbtag '100'
RP/0/RSP0/CPU0:router(config-rp-gl)# end-global
```

Related Commands

Command	Description
policy-global, on page 1712	Enters global parameter configuration mode.

end-policy

To end the definition of a route policy and exit route-policy configuration mode, use the **end-policy** command in route-policy configuration mode.

end-policy

Syntax Description

This command has no arguments or keywords.

Command Default

No default behavior or values

Command Modes

Route-policy 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **end-policy** command to end the definition of a route policy and exit route-policy configuration mode.

Task ID

Task ID	Operations
route-policy	read, write

Examples

In the following example, the **end-policy** command ends the definition of a route policy:

```
RP/0/RSP0/CPU0:router(config)#route-policy med-to-local-pref
RP/0/RSP0/CPU0:router(config-rpl)#if med eq 150 then
RP/0/RSP0/CPU0:router(config-rpl-if)# set local-preference 10
RP/0/RSP0/CPU0:router(config-rpl-if)# elseif med eq 200 then
RP/0/RSP0/CPU0:router(config-elseif)# set local-preference 60
RP/0/RSP0/CPU0:router(config-elseif)# elseif med eq 250 then
RP/0/RSP0/CPU0:router(config-elseif)# set local-preference 0

RP/0/RSP0/CPU0:router(config-elseif)# endif
RP/0/RSP0/CPU0:router(config-rpl)# end-policy
```

Related Commands

Command	Description
route-policy (RPL), on page 1729	Defines a route policy and enters route-policy configuration mode.

end-set

To end the definition of an AS path set, a prefix set, a community set, an extended community set, or an RD set and return to global configuration mode, use the **end-set** command in route-policy configuration mode.

end-set

Syntax Description

This command has no arguments or keywords.

Command Default

No default behavior or values

Command Modes

AS path set configuration

Prefix set configuration

Community set configuration

Extended community set configuration

Route distinguisher set 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **end-set** command to end the definition of an AS path set, a prefix set, a community set, or an extended community set.

Task ID

Task ID	Operations
route-policy	read, write

Examples

In the following example, the **end-set** command ends the definition of an AS path set named aset1:

```
RP/0/RSP0/CPU0:router(config)# as-path-set aset1
RP/0/RSP0/CPU0:router(config-as)# ios-regex '_42$',
RP/0/RSP0/CPU0:router(config-as)# ios-regex '_127$'

RP/0/RSP0/CPU0:router(config-as)# end-set
RP/0/RSP0/CPU0:router(config)#
```

The following example shows how to create an RD set called my_rd_set and use the **end-set** command to end the definition:

```
RP/0/RSP0/CPU0:router(config)# rd-set my_rd_set
```

```
RP/0/RSP0/CPU0:router(config-rd)# 172.16.0.0/16:*,
RP/0/RSP0/CPU0:router(config-rd)# 172.17.0.0/16:100,
RP/0/RSP0/CPU0:router(config-rd)# 192:*,
RP/0/RSP0/CPU0:router(config-rd)# 192:100
RP/0/RSP0/CPU0:router(config-rd)# end-set
```

event prefix route-policy pass

To get additional trace information about a particular set of prefixes, use the **event prefix route-policy pass** command in global address family configuration mode.

event prefix route-policy pass

Syntax Description

This command has no arguments or keywords.

Command Default

No default behavior or values

Command Modes

Global address family

Command History

Release	Modification
Release 7.3.1	This command was introduced.

Usage Guidelines

- Use this command strictly for debugging purposes when general BGP traces are insufficient.
- Enabling this command increases the volume of trace output, providing more detailed information for a specific set of prefixes.
- **Do not enable this command to trace all prefixes**, as using a pass-all route policy will degrade system performance and reduce the clarity of information for the prefixes you actually need to debug.
- Apply this command selectively to a limited subset of prefixes to avoid overwhelming the system and to maintain useful trace data.
- Remove or disable the command once debugging is complete to prevent ongoing performance impact.

Example

The following example shows how to gather additional trace information about a particular set of prefixes:

```
Router(config)# prefix-set prefix
Router(config-pfx)# 10.2.0.0/16
Router(config-pfx)# exit
Router(config)#
Router(config)# route-policy tracing
Router(config-rpl)# if destination in prefix then pass else drop endif
Router(config-rpl)# exit
Router(config-rpl)# router bgp 1
Router(config-bgp)# address-family ipv4  unicast
Router(config-bgp-af)# event prefix route-policy

/*If there are any changes related to the prefix
10.2.0.0/16, the same changes will be seen when
using the show bgp trace per-prefix command. */
```

extcommunity rt is-empty

To check if a Border Gateway Protocol (BGP) route has route target (RT) extended community attributes associated with it, use the **extcommunity rt is-empty** command in route-policy configuration mode.

extcommunity rt is-empty

Syntax Description

This command has no arguments or keywords.

Command Default

No default behavior or value

Command Modes

Route-policy 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **extcommunity rt is-empty** command as a conditional expression within an **if** statement to check if a BGP route has extended community attributes associated with it.

Note For a list of all conditional expressions available within an **if** statement, see the **if** command.

The **is-empty** operator takes no arguments and evaluates to true if the route has no extended community attributes associated with it.

Task ID

Task ID	Operations
route-policy	read, write

Examples

In the following example, if the extended community is empty, then the local preference is set to 100:

```
RP/0/RSP0/CPU0:router(config)# route-policy extcommunity-is-empty-example
RP/0/RSP0/CPU0:router(config-rpl)# if extcommunity rt is-empty then
RP/0/RSP0/CPU0:router(config-rpl-if)# set local-preference 100

RP/0/RSP0/CPU0:router(config-rpl-if)# endif
RP/0/RSP0/CPU0:router(config-rpl)# end-policy
```

extcommunity rt matches-any

To match any element of a Border Gateway Protocol (BGP) route target (RT) extended community set, use the **extcommunity rt matches-any** command in route-policy configuration mode.

extcommunity rt matches-any {*extcommunity-set-nameinline-extcommunity-setparameter*}

Syntax Description

extcommunity-set-name	Name of an RT extended community set.
inline-extcommunity-set	Inline RT extended community set. The inline extended community set must be enclosed in parentheses.
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 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **extcommunity rt matches-any** command as a conditional expression within an **if** statement to match elements of an extended community set.



Note For a list of all conditional expressions available within an **if** statement, see the **if** command.

A simple condition using the **matches-any** operator evaluates as true if at least one extended community in the route matches an extended community specification in the named or inline set. If no extended community in the route matches any of the specifications in the named or inline set, then this simple condition evaluates to false. Likewise, when there is no extended community at all in the route, the condition evaluates to false.

Matching an extended community in the route to a specification in a named or an inline set is intuitive. In inline sets, extended community specifications may be parameterized, in which case the relevant matching is done when the value of the parameter has been supplied.

Task ID

Task ID	Operations
route-policy	read, write

Examples

In the following example, an extended community set named my-extcommunity-set and a parameterized route-policy named my-extcommunity-set-example($tag,$ip) are defined. The **extcommunity rt matches-any** command is used in an if statement such that if at least one extended community in the route matches an extended community specification in the named set, then the local preference is set to 100. If there is no extended community in the route that matches any of the specifications in the named set, then the condition evaluates as false and the extended community is compared to the inline extended sets.

```
RP/0/RSP0/CPU0:router(config)# extcommunity-set rt my-extcommunity-set
RP/0/RSP0/CPU0:router(config-ext)# 10:615,
RP/0/RSP0/CPU0:router(config-ext)# 10:6150,
RP/0/RSP0/CPU0:router(config-ext)# 15.15.15.15:15
RP/0/RSP0/CPU0:router(config-ext)# end-set

RP/0/RSP0/CPU0:router(config)# route-policy my-extcommunity-set-example($tag,$ip)
RP/0/RSP0/CPU0:router(config-rpl)# if extcommunity rt matches-any my-extcommunity-set then
RP/0/RSP0/CPU0:router(config-rpl-if)# set local-preference 100
RP/0/RSP0/CPU0:router(config-rpl-if)# elseif extcommunity rt matches-any (10:20, 10:$tag)
then
RP/0/RSP0/CPU0:router(config-rpl-elseif)# set local-preference 200
RP/0/RSP0/CPU0:router(config-rpl-elseif)# elseif extcommunity rt matches-any ($ip:$tag)
then
RP/0/RSP0/CPU0:router(config-rpl-elseif)# set local-preference 300
RP/0/RSP0/CPU0:router(config-rpl-elseif)# elseif extcommunity rt matches-any (2.3.4.5:$tag)
 then
RP/0/RSP0/CPU0:router(config-rpl-elseif)# set local-preference 400
RP/0/RSP0/CPU0:router(config-rpl-elseif)# endif
RP/0/RSP0/CPU0:router(config-rpl)# end-policy
```

Related Commands

Command	Description
extcommunity rt matches-every, on page 1657	Matches every element of a BGP RT extended community set.

extcommunity rt matches-every

To match every element of a Border Gateway Protocol (BGP) route target (RT) extended community set, use the **extcommunity rt matches-every** command in route-policy configuration mode.

extcommunity rt matches-every {*extcommunity-set-nameinline-extcommunity-setparameter*}

Syntax Description

extcommunity-set-name	Name of an RT extended community set.
inline-extcommunity-set	Inline RT extended community set. The inline extended community set must be enclosed in parentheses.
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 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **extcommunity rt matches-every** command as a conditional expression within an **if** statement to match every element of an RT extended community set.



Note For a list of all conditional expressions available within an **if** statement, see the **if** command.

A simple condition using the **matches-every** operator evaluates as true if every extended community value in the extended community attribute for the route matches at least one element of the extended community set or inline set. If no extended community in the route matches any of the specifications in the named or inline set, then this simple condition evaluates to false. Likewise, when there is no extended community at all in the route, the condition evaluates to false.

Matching an extended community in the route to a specification in a named or an inline set is intuitive. In inline sets, extended community specifications may be parameterized, in which case the relevant matching is done when the value of the parameter has been supplied.

Task ID

Task ID	Operations
route-policy	read, write

Examples

In the following example, an extended community set named my-extcommunity-set and a parameterized route policy named extcommunity-matches-every-example ($as, $tag) are defined. The condition extcommunity rt matches-every is used in an if statement in this policy. If it evaluates to true, the local-preference value is set to 100. If it evaluates to false, the extended community is evaluated using an inline set. If that condition evaluates to true, the local-preference value is set to 200. If it evaluates to false, the local-preference value is set to 300.

```
RP/0/RSP0/CPU0:router(config)# extcommunity-set rt my-extcommunity-set
RP/0/RSP0/CPU0:router(config-ext)# 10:20,
RP/0/RSP0/CPU0:router(config-ext)# 10:30,
RP/0/RSP0/CPU0:router(config-ext)# 10:40
RP/0/RSP0/CPU0:router(config-ext)# end-set
RP/0/RSP0/CPU0:router(config)# route-policy extcommunity-matches-every-example($as,$tag)
RP/0/RSP0/CPU0:router(config-rpl)# if extcommunity rt matches-every my-extcommunity-set
then
RP/0/RSP0/CPU0:router(config-rpl-if)# set local-preference 100
RP/0/RSP0/CPU0:router(config-rpl-if)# elseif extcommunity rt matches-every (10:20, 10:$tag,
 $as:30) then
RP/0/RSP0/CPU0:router(config-rpl-elseif)# set local-preference 200
RP/0/RSP0/CPU0:router(config-rpl-elseif)# elseif
RP/0/RSP0/CPU0:router(config-rpl-elseif)# set local-preference 300
RP/0/RSP0/CPU0:router(config-rpl-elseif)# endif
RP/0/RSP0/CPU0:router(config-rpl)# end-policy
```

Related Commands

Command	Description
extcommunity rt matches-any, on page 1655	Matches any element of a BGP RT extended community set.

extcommunity rt matches-within

To match at least one element of an extended community set of a Border Gateway Protocol (BGP) route target (RT), use the **extcommunity rt matches-within** command in route-policy configuration mode.

extcommunity rt matches-within { *rt-type-extcommunity-set-name inline-extcommunity-set parameter* }

Syntax Description

rt-type-extcommunity-set-name	Name of an RT extended community set.
inline-extcommunity-set	Inline RT extended community set, enclosed in parentheses.
parameter	Parameter name preceded with a “$” symbol.

Command Default

None

Command Modes

Route-policy configuration

Command History

Release	Modification
Release 4.2.0	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.

Use the extcommunity rt matches-within command as a conditional expression within an if statement to match elements of an extended community set.

Note For a list of all conditional expressions available within an if statement, see the if command.

A simple condition using the matches-within operator evaluates as true if all the elements in extended community from the route match any element in the extended community set. For example, let 'c' be the RTs from the route and 'm' be the RT set from the policy. With the **extcommunity rt matches-within** configuration, each value in 'c' must match any (or at least one) value in 'm'.

Matching an extended community in the route to a specification in a named or an inline set is intuitive. In inline sets, extended community specifications may be parameterized, in which case the relevant matching is done when the value of the parameter has been supplied.

Task ID

Task ID	Operation
route-policy	read, write

In the following example, an extended community set named *my-extcommunity-set* and a parameterized route-policy named *my-extcommunity-set-example($tag,$ip)* are defined. The **extcommunity rt matches-within** command is used in an if statement such that if all extended community values in the route match any element of the extended community specification in the named set, then the local preference is set to 100.

```
Router(config)#extcommunity-set rt my-extcommunity-set
Router(config-ext)#10:615,
Router(config-ext)#10:6150,
Router(config-ext)#15.15.15.15:15
Router(config-ext)#end-set
Router(config)#route-policy my-extcommunity-set-example($tag,$ip)
Router(config-rpl)#if extcommunity rt matches-within my-extcommunity-set then
Router(config-rpl-if)#set local-preference 100
```

extcommunity-set bandwidth

To define a Border Gateway Protocol (BGP) route target (RT) extended community set, use the **extcommunity-set bandwidth** command in global configuration mode. To remove the RT community set, use the **no** form of this command.

extcommunity-set bandwidth *name*
no extcommunity-set bandwidth *name*

Syntax Description

name	Name of bandwidth extended community set.

Command Default

No default behavior or values

Command Modes

Global 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

The **extcommunity-set bandwidth** extended community set advertises the bandwidth of an autonomous system exit link as an extended community for links between directly connected external BGP (eBGP) neighbors. The link bandwidth extended community attribute is propagated to iBGP peers when extended community exchange is enabled. The extended community is used with BGP multipath features to configure load balancing over links with unequal bandwidth. When this community set is enabled, the routes learned from the external neighbors are propagated through the internal BGP (iBGP) network with the bandwidth of the source external link.

The bandwidth extended community supports the following option:

- *X:Y—X* indicates the ASN of the router sending the extended community. This ASN is a 16-bit decimal number. *Y* indicates the bandwidth in bytes per second. The bandwidth is a 32-bit decimal number.

Examples

In the following example, a bandwidth extended community set named `extcomm-bw` is defined:

```
RP/0/RSP0/CPU0:router(config)# extcommunity-set bandwidth extcomm-bw
RP/0/RSP0/CPU0:router(config-ext)# 10002:777
RP/0/RSP0/CPU0:router(config-ext)# end-set
```

extcommunity-set cost

To define a cost extended community set, use the **extcommunity-set cost** command in global configuration mode. To remove the cost extended community set, use the **no** form of this command.

extcommunity-set cost *name*
no extcommunity-set cost *name*

Syntax Description

name	Name of a cost extended community set. The *name* argument is case sensitive, can contain any alphanumeric characters, and can be up to 63 characters in length.

Command Default

No default behavior or values

Command Modes

global configuration

Command History

Release	Modification
Release 3.7.2	This command was introduced.
Release 3.9.0	Support was added for more cost extended community formats.

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.

Use the **extcommunity-set cost** command to define a cost extended community set.

An extended community set is analogous to a community set except that it contains extended community values instead of regular community values. Extended community values are 64-bit structured values. An extended community set also supports named forms and inline forms.

Cost extended communities can be entered in these formats:

- *0-255*—Decimal number
- **igp:**—Cost Community with IGP as point of insertion
- **pre-bestpath:**—Cost Community with Pre-Bestpath as point of insertion

Multiple cost community set clauses can be configured in each route policy block or sequence. Each cost community set clause must have a different ID (0-255). The cost community set clause with the lowest cost-value is preferred by the best path selection process when all other attributes are equal.

As with community sets, the inline form supports parameterization within parameterized policies. Either portion of the extended community value can be parameterized.

Every extended community set must contain at least one extended community value. Empty extended community sets are invalid and the policy configuration system rejects them.

Wildcards (*) and regular expressions are allowed for extended community set elements.

Examples

In the following example, a cost extended community set named `extcomm-cost` is defined:

```
RP/0/RSP0/CPU0:router(config)# extcommunity-set cost extcomm-cost
RP/0/RSP0/CPU0:router(config-ext)# IGP:90:914,
RP/0/RSP0/CPU0:router(config-ext)# Pre-Bestpath:91:915
RP/0/RSP0/CPU0:router(config-ext)# end-set
```

extcommunity-set opaque

To define a Border Gateway Protocol (BGP) Opaque extended community set, use the **extcommunity-set opaque** command in global configuration mode. To remove the Opaque extended community set, use the **no** form of this command.

extcommunity-set opaque *name*
no extcommunity-set opaque *name*

Syntax Description

name	Name of color extended community.

Command Default

No default behavior or values

Command Modes

Global configuration

Command History

Release	Modification
Release 6.2.1	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.

Use the **extcommunity-set opaque** command to define an Opaque extended community set that contains color extended communities.

Opaque extended communities support these options:

- *****—Wildcard (any community or part thereof)
- *1-4294967295*—color value

Examples

In the following example, an Opaque extended community set named `color78` is defined for BGP extended communities. When RPL `a_color` is applied as import route-policy to VRF, the route with only color 7 is dropped while the route with both color 7 and color 8 is passed:

```
RP/0/RSP0/CPU0:router(config)# extcommunity-set opaque color78
RP/0/RSP0/CPU0:router(config-ext)# 7,
RP/0/RSP0/CPU0:router(config-ext)# 8
RP/0/RSP0/CPU0:router(config-ext)# end-set

RP/0/RSP0/CPU0:router(config)# route-policy a_color
RP/0/RSP0/CPU0:router(config-rpl)# if extcommunity color matches-every mycolor78 then
RP/0/RSP0/CPU0:router(config-rpl-if)# set extcommunity color color27128
RP/0/RSP0/CPU0:router(config-rpl-if)# endif
RP/0/RSP0/CPU0:router(config-rpl)# end-policy
```

extcommunity-set rt

To define a Border Gateway Protocol (BGP) route target (RT) extended community set, use the **extcommunity-set rt** command in global configuration mode. To remove the RT community set, use the **no** form of this command.

extcommunity-set rt *name*
no extcommunity-set rt *name*

Syntax Description

name	Name of an RT extended community set.

Command Default

No default behavior or values

Command Modes

Global configuration

Command History

Release	Modification
Release 3.7.2	This command was introduced.
Release 3.9.0	Support was added for more rt extended community formats.

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.

Use the **extcommunity-set rt** command to define an RT extended community set for BGP.

Regular expressions and ranges can be specified to match the extended communities. Regular expressions and ranges can be specified in an extended community set to support the matching of communities. An attempt to use an extended community set that contains a range or regular expression to set an extended community set value is rejected when an attempt to attach such a policy is made.

An extcommunity set RT holds RT extended community values to match against the Border Gateway Protocol (BGP) RT extended community attribute. RT extended communities can be entered in these formats:

- *****—Wildcard (any community or part thereof)
- *1-4294967295*—32-bit decimal number
- *1-65535*—16-bit decimal number
- *A.B.C.D/M:N* —Extended community - IPv4 prefix format
- *A.B.C.D:N*—Extended community - IPv4 format
- *ASN:N*—Extended community - ASPLAIN format
- *X.Y:N*—Extended community - ASDOT format
- **dfa-regex**—DFA (deterministic finite automata) style regular expression
- **ios-regex**—Traditional IOS style regular expression
- **match**—Match option for the set. Default - any

Note The `dfa-regex` and `ios-regex` syntax for community set is *"[']*[^':&<>]*:[^':&<>]*[']"*. This means that regex starts with a single-quote (') followed by a string of any character (that does not include single-quote, colon, ampersand, less-than, greater-than, or space) followed by a colon, and a string of any characters (that does not include single-quote, colon, ampersand, less-than, greater-than, or space) followed by single-quote.

N is a number within the range of 1 to 65535.

Examples

In the following example, an RT extended community set named `extcomm-rt` is defined:

```
RP/0/RSP0/CPU0:router(config)# extcommunity-set rt extcomm-rt
RP/0/RSP0/CPU0:router(config-ext)# 10002:666
RP/0/RSP0/CPU0:router(config-ext)# 10.0.0.2:666
RP/0/RSP0/CPU0:router(config-ext)# end-set
```

extcommunity-set seg-nh

To define a Point-to-Multipoint (P2MP) segmented next-hop extended community use the **extcommunity-set seg-nh** command in global configuration mode. To remove the seg-nh extended community set, use the **no** form of this command.

extcommunity-set seg-nh *name*
no extcommunity-set seg-nh *name*

Syntax Description

name	Indicates the Point-to-Multipoint (P2MP) segmented next-hop extended community name to indicate segmented next hop.

Command Default

No default behavior or values

Command Modes

Global 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Segmented next hop extended community supports the following option:

- *A.B.C.D*—Extended community - IPv4 address format

Examples

In the following example, a segmented next hop extended community set named `extcomm-nh` is defined for BGP extended communities.

```
RP/0/RSP0/CPU0:router(config)# extcommunity-set seg-nh extcomm-nh
RP/0/RSP0/CPU0:router(config-ext)# 1.1.1.1
RP/0/RSP0/CPU0:router(config-ext)# end-set
```

extcommunity-set soo

To define a Border Gateway Protocol (BGP) Site-of-Origin (SoO) extended community set, use the **extcommunity-set soo** command in global configuration mode. To remove the SoO extended community set, use the **no** form of this command.

extcommunity-set soo *name*
no extcommunity-set soo *name*

Syntax Description

name	Name of an SoO extended community set.

Command Default

No default behavior or values

Command Modes

Global configuration

Command History

Release	Modification
Release 3.7.2	This command was introduced.
Release 3.9.0	Support was added for more soo extended community formats.

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.

Use the **extcommunity-set soo** command to define an SoO extended community set.

An extcommunity set soo holds SoO extended community values to match against the Border Gateway Protocol (BGP) SoO extended community attribute. SoO extended communities can be entered in these formats:

- *****—Wildcard (any community or part thereof)
- *1-4294967295*—32-bit decimal number
- *1-65535*—16-bit decimal number
- *A.B.C.D/M:N*—Extended community - IPv4 prefix format
- *A.B.C.D:N*—Extended community - IPv4 format
- *ASN:N*—Extended community - ASPLAIN format
- *X.Y:N*—Extended community - ASDOT format
- **dfa-regex**—DFA style regular expression
- **ios-regex**—Traditional IOS style regular expression
- **match** —Match option for the set. Default - any

N is a site-specific number.

Examples

In the following example, a SoO extended community set named `extcomm-soo` is defined:

```
RP/0/RSP0/CPU0:router(config)# extcommunity-set soo extcomm-soo
RP/0/RSP0/CPU0:router(config-ext)# 66:60001,
RP/0/RSP0/CPU0:router(config-ext)# 77:70001,
RP/0/RSP0/CPU0:router(config-ext)# 88:80001,
RP/0/RSP0/CPU0:router(config-ext)# 99:90001,

RP/0/RSP0/CPU0:router(config-ext)# 100.100.100.1:153
RP/0/RSP0/CPU0:router(config-ext)# end-set
```

extcommunity soo is-empty

To determine if a Border Gateway Protocol (BGP) route has any Site-of-Origin (SoO) extended communities associated with it, use the **extcommunity soo is-empty** command in route-policy configuration mode.

extcommunity soo is-empty

Syntax Description

This command has no arguments or keywords.

Command Default

No default behavior or values

Command Modes

Route-policy 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **extcommunity soo is-empty** command as a conditional expression within an **if** statement to check if a BGP SoO route has extended community attributes associated with it.

Note For a list of all conditional expressions available within an **if** statement, see the **if** command.

The **is-empty** operator takes no arguments and evaluates to true if the route has no SoO extended community attributes associated with it.

Task ID

Task ID	Operations
route-policy	read, write

Examples

In the following example, if a route has no SoO extended communities associated with it, the local preference is set to 100:

```
RP/0/RSP0/CPU0:router(config)# route-policy extcommunity-is-empty-example
RP/0/RSP0/CPU0:router(config-rpl)# if extcommunity soo is-empty then
RP/0/RSP0/CPU0:router(config-rpl-if)# set local-preference 100
RP/0/RSP0/CPU0:router(config-rpl-if)# endif
RP/0/RSP0/CPU0:router(config-rpl)# end-policy
```

extcommunity soo matches-any

To match any element of a Border Gateway Protocol (BGP) Site-of-Origin (SoO) extended community set, use the **extcommunity soo matches-any** command in route-policy configuration mode.

extcommunity soo matches-any {*extcommunity-set-nameinline-extcommunity-setparameter*}

Syntax Description

extcommunity-set-name	Name of a SoO extended community set.
inline-extcommunity-set	Inline SoO extended community set. The inline extended community set must be enclosed in parentheses.
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 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **extcommunity soo matches-any** command as a conditional expression within an **if** statement to match elements of an extended community set.

Note For a list of all conditional expressions available within an **if** statement, see the **if** command.

A simple condition using the **matches-any** operator evaluates as true if at least one extended community in the route matches an extended community specification in the named or inline set. If no extended community in the route matches any of the specifications in the named or inline set, then this simple condition evaluates to false. Likewise, when there is no extended community at all in the route, the condition evaluates to false.

Matching an extended community in the route to a specification in a named or an inline set is intuitive. In inline sets, extended community specifications may be parameterized, in which case the relevant matching is done when the value of the parameter has been supplied.

Task ID

Task ID	Operations
route-policy	read, write

Examples

In the following example, an SoO extended community set named extcomm-soo and a parameterized route policy named my-extcommunity-set-example($tag,$ip) are defined.

The condition route policy named extcommunity soo matches-any is used in an if statement in this policy. If it evaluates to true, the local preference value is set to 100.

If it evaluates to false, the SoO extended community is evaluated using an inline set. If it evaluates to true, the local preference value is set to 200.

If it evaluates to false, the SoO extended community is evaluated using a different inline set. If it evaluates to true, the local preference value is set to 300.

If it evaluates to false, the SoO extended community is evaluated using a different inline set. If it evaluates to true, the local preference value is set to 400.

```
RP/0/RSP0/CPU0:router(config)# extcommunity-set soo extcomm-soo
RP/0/RSP0/CPU0:router(config-ext)# 66:60001,
RP/0/RSP0/CPU0:router(config-ext)# 77:70001,
RP/0/RSP0/CPU0:router(config-ext)# 88:80001,
RP/0/RSP0/CPU0:router(config-ext)# 99:90001,
RP/0/RSP0/CPU0:router(config-ext)# 100.100.100.1:153
RP/0/RSP0/CPU0:router(config-ext)# end-set

RP/0/RSP0/CPU0:router(config)# route-policy my-extcommunity-set-example($tag,$ip)
RP/0/RSP0/CPU0:router(config-rpl)# if extcommunity soo matches-any extcomm-soo then
RP/0/RSP0/CPU0:router(config-rpl-if)# set local-preference 100
RP/0/RSP0/CPU0:router(config-rpl-if)# elseif extcommunity soo matches-any (10:20, 10:$tag)
 then
RP/0/RSP0/CPU0:router(config-rpl-elseif)# set local-preference 200
RP/0/RSP0/CPU0:router(config-rpl-elseif)# elseif extcommunity soo matches-any ($ip:$tag)
then
RP/0/RSP0/CPU0:router(config-rpl-elseif)# set local-preference 300
RP/0/RSP0/CPU0:router(config-rpl-elseif)# elseif extcommunity soo matches-any (2.3.4.5:$tag)
 then
RP/0/RSP0/CPU0:router(config-rpl-elseif)# set local-preference 400
RP/0/RSP0/CPU0:router(config-rpl-elseif)# endif
RP/0/RSP0/CPU0:router(config-rpl)# end-policy
```

Related Commands

Command	Description
extcommunity rt matches-any, on page 1655	Matches any element of a BGP RT extended community set.
extcommunity soo matches-every, on page 1673	Matches every element of a BGP SoO extended community set.

extcommunity soo matches-every

To match every element of a Border Gateway Protocol (BGP) Site-of-Origin (SoO) extended community set, use the **extcommunity soo matches-every** command in route-policy configuration mode.

extcommunity soo matches-every {*extcommunity-set-nameinline-extcommunity-setparameter*}

Syntax Description

extcommunity-set-name	Name of a SoO extended community set.
inline-extcommunity-set	Inline SoO extended community set. The inline extended community set must be enclosed in parentheses.
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 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **extcommunity soo matches-every** command as a conditional expression within an **if** statement to match every element of a SoO extended community set.

Note For a list of all conditional expressions available within an **if** statement, see the **if** command.

A simple condition using the **matches-every** operator evaluates as true if every extended community value in the extended community attribute for the route matches at least one element of the extended community set or inline set. If no extended community in the route matches any of the specifications in the named or inline set, then this simple condition evaluates to false. Likewise, when there is no extended community at all in the route, the condition evaluates to false.

Matching an extended community in the route to a specification in a named or an inline set is intuitive. In inline sets, extended community specifications may be parameterized, in which case the relevant matching is done when the value of the parameter has been supplied.

Examples

In the following example, an extended community set named my-extcomm-rt-set and a parameterized route policy named extcommunity-matches-every-example($as, $tag) are defined. The condition extcommunity soo matches-every is used in an if statement in this policy and if it evaluates to true, the local-preference value is set to 100. If it evaluates to false, the extended community is evaluated using an inline set. If that condition evaluates to true, the local-preference value is set to 200. If it evaluates to false, the local-preference value is set to 300.

```
RP/0/RSP0/CPU0:router(config)# extcommunity-set soo my-extcomm-rt-set
RP/0/RSP0/CPU0:router(config-ext)# 10:20,
RP/0/RSP0/CPU0:router(config-ext)# 10:30,
RP/0/RSP0/CPU0:router(config-ext)# 10:40
RP/0/RSP0/CPU0:router(config-ext)# end-set

RP/0/RSP0/CPU0:router(config)# route-policy extcommunity-matches-every-example($as, $tag)
RP/0/RSP0/CPU0:router(config-rpl)# if extcommunity soo matches-every my-extcomm-rt-set then
RP/0/RSP0/CPU0:router(config-rpl-if)# set local-preference 100
RP/0/RSP0/CPU0:router(config-rpl-if)# elseif extcommunity soo matches-every (10:20, 10:$tag,
 $as:30) then
RP/0/RSP0/CPU0:router(config-rpl-elseif)# set local-preference 200
RP/0/RSP0/CPU0:router(config-rpl-elseif)# else
RP/0/RSP0/CPU0:router(config-rpl-elseif)# set local-preference 300
RP/0/RSP0/CPU0:router(config-rpl-elseif)# endif
RP/0/RSP0/CPU0:router(config-rpl)# end-policy
```

Related Commands

Command	Description
extcommunity soo matches-any, on page 1671	Matches any element of a BGP SoO extended community set.

globalVar*N* is

To check the value of globalVar*N* value assigned through the **var globalVar***N* command, use the **globalVar***N* **is** command in router-policy configuration mode.

globalVar*N* **is** *{number | parameter}*

Syntax Description

number	Value assigned to a 32-bit unsigned integer. Range is from 1 to 4294967295.
parameter	Parameter name. The parameter name must be preceded with a "$."

Command Default

If the globalVar*N* is not assigned using the **var globalVar***N* command, then the default value for globalVar*N* is zero.

Command Modes

Route-policy configuration

Command History

Release	Modification
Release 5.1.3	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.

The scope of this variable is from the start of policy to end of policy attached under any attach-point. The variable can be assigned in any child policy and can check for value in parent policy or any other hierarchy of route-policy referred using apply statement.

Task ID

Task ID	Operation
route-policy	read, write

Example

This example is to identify routes which has communities configured under community-sets internal_set and service2 and set different local-preference values based on the community values.

```
community-set internal_set
         1:1
end-set
community-set service2
         3:3
end-set
route-policy child1
 var globalVar1 100
 #... user can add more actions or conditions...
end-policy
route-policy child2
 var globalVar2 200
```

```
 #...user can add more actions or conditions ..
end-policy
route-policy inbound
  if community matches-any internal_set then
    apply child1
  endif
  if community matches-any service2 then
    apply child2
  endif
  if globalVar1 is 100 and globalVar2 is 200 then
     set local-preference 250
  elseif globalVar1 is 100 then
     set local-preference 150
  elseif globalVar2 is 200 then
     set local-preference 50
  endif
end-policy
```

if

To decide which actions or dispositions should be taken for a given route, use the **if** command in route-policy configuration mode.

if *conditional-expression* **then** *action-statement* [*action-statement*] [**elseif** *conditional-expression* **then** *action-statement* [*action-statement*]] [**else** *action-statement* [*action-statement*]] **endif**

Syntax Description

conditional-expression	Expression to decide which actions or dispositions should be taken for the given route.
then	Executes an action statement if the **if** condition is true.
elseif	Strings together a sequence of tests.
else	Executes an action statement if the **if** condition is false.
endif	Ends the **if** statement.
action-statement	Sequence of operations that modify a route.

Command Default

No default behavior or values

Command Modes

Route-policy 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

The **if** command uses a conditional expression to decide which actions or dispositions should be taken for a given route. Table 173: Conditional Expressions, on page 1677 lists the conditional expressions.

An action statement is a sequence of operations that modify a route, most of which are distinguished by the **set** keyword. In a route policy, these operations can be grouped. Table 174: Action Statements, on page 1679 lists the action statements.

Table 173: Conditional Expressions

Command	Description
as-path in, on page 1605	Matches the AS path of a route to an AS path set. The AS path is a sequence of autonomous system numbers traversed by a route.
as-path is-local, on page 1607	Determines if the router (or another router within this autonomous system or confederation) originated the route.

Command	Description
as-path length, on page 1609	Performs a conditional check based on the length of the AS path.
as-path neighbor-is, on page 1611	Tests the autonomous system number or numbers at the head of the AS path against a sequence of one or more integral values or parameters.
as-path originates-from, on page 1614	Tests an AS path against the AS sequence beginning with the AS number that originated a route.
as-path passes-through, on page 1616	Tests to learn if the specified integer or parameter appears anywhere in the AS path or if the sequence of integers and parameters appears.
as-path unique-length, on page 1622	Performs specific checks based on the length of the AS path.
community is-empty, on page 1624	Learns if a route has community attributes associated with it.
community matches-any, on page 1625	Matches any element of a community set.
community matches-every, on page 1627	Matches every element of a community set.
destination in, on page 1640	Matches a destination entry in a named prefix set or inline prefix set.
extcommunity rt is-empty, on page 1654	Learns if a route has RT extended community attributes associated with it.
extcommunity rt matches-any, on page 1655	Matches elements of an RT extended community set.
extcommunity rt matches-every, on page 1657	Matches every element of an RT extended community set.
extcommunity soo is-empty, on page 1670	Learns if a route has SoO extended community attributes associated with it.
extcommunity soo matches-any, on page 1671	Matches elements of an SoO extended community set.
extcommunity soo matches-every, on page 1673	Matches every element of an SoO extended community set.
local-preference, on page 1687	Specifies BGP local-preference attribute
med, on page 1697	Compares the MED to an integer value or a parameterized value.
next-hop in, on page 1699	Compares the next-hop associated with the route to data contained in either a named or an inline prefix set.
orf prefix in, on page 1701	Matches a prefix in a prefix set or an inline prefix set.
origin is, on page 1703	Tests the value of the origin attribute.
path-type is, on page 1711	Tests the path type.

Command	Description
protocol, on page 1718	Checks if a protocol is installing the route.
rd in, on page 1720	Compares the RD associated with the route to data contained in either a named or an inline RD set.
rib-has-route, on page 1726	Checks if a route is in the RIB.
route-has-label, on page 1728	Checks if a route has a Multiprotocol Label Switching (MPLS) label.
route-type is, on page 1731	Compares route types when redistribution is being performed into BGP, OSPF, or IS-IS.
source in, on page 1886	Tests the source of the route against the data in either a named or an inline prefix set.
tag, on page 1890	Matches a specific tag value.
tag in, on page 1891	Conditionally compares tag-route against tag-set.
vpn-distinguisher is, on page 1897	Compares the VPN distinguisher against a specified value.

Table 174: Action Statements

Command	Description
abort (RPL), on page 1599	Discards a route policy definition and returns to global configuration mode.
add, on page 1601	Adds an offset to an existing value.
apply, on page 1603	Executes a parameterized or an unparameterized policy from within another policy.
delete community, on page 1634	Deletes community values from a community list in a route.
delete extcommunity rt, on page 1638	Deletes extended community values from an extended community list in a route.
done, on page 1642	Accepts this route with no further processing
drop, on page 1644	Drops a route.
end-policy, on page 1650	Ends the definition of a route policy and exits route-policy configuration mode.
pass, on page 1709	Signifies that even though the route has not been modified, the user wants to continue executing in the policy block.
prepend as-path, on page 1716	Prepends the AS path with additional autonomous system numbers.
replace as-path, on page 1723	Replaces a sequence of AS numbers or private AS numbers in the AS path with the configured local AS.

Command	Description
set community, on page 1740	Sets the BGP community attribute.
set dampening, on page 1743	Configures BGP route dampening.
set eigrp-metric, on page 1746	Sets the Enhanced Interior Gateway Routing Protocol (EIGRP) metric value.
set extcommunity cost, on page 1749	Replaces or adds the extended communities for a cost on the route.
set extcommunity rt, on page 1751	Replaces or adds the extended communities for an RT on the route.
set ip-precedence, on page 1753	Sets the IP precedence to classify packets.
set isis-metric, on page 1755	Sets the IS-IS metric attribute value.
set label, on page 1756	Sets the BGP label attribute value.
set level, on page 1761	Configures the IS-IS level in which redistributed routes should be sent.
set local-preference, on page 1762	Specifies a preference value for the autonomous system path.
set med, on page 1763	Sets the MED value.
set metric-type (IS-IS), on page 1765	Controls whether IS-IS treats the metric as an internal or external metric.
set metric-type (OSPF), on page 1766	Controls whether OSPF treats the cost as a Type 1 or Type 2 metric.
set next-hop, on page 1767	Replaces the next-hop associated with a given route.
set origin, on page 1769	Changes the origin attribute.
set ospf-metric, on page 1770	Sets an OSPF protocol metric attribute value.
set qos-group (RPL), on page 1773	Sets the QoS group to classify packets.
set rib-metric, on page 1774	Sets a RIB metric attribute value for a table policy.
set rip-metric, on page 1775	Sets RIP metric attributes.
set rip-tag, on page 1776	Sets route tag attribute.
set tag, on page 1781	Sets the tag attribute.
set traffic-index, on page 1782	Sets the traffic index attribute.
set weight, on page 1785	Sets the weight value for BGP routes.
suppress-route, on page 1889	Indicates that a given component of an aggregate should be suppressed, that is, not advertised.
unsuppress-route, on page 1894	Indicates that a given component of an aggregate should be unsuppressed.

Command	Description
set vpn-distinguisher, on page 1784	Sets the VPN distinguisher value.

Task ID

Task ID	Operations
route-policy	read, write

Examples

In the following example, any route whose AS path is in the set as-path-set-1 is dropped:

```
RP/0/RSP0/CPU0:router(config-rpl)# if as-path in as-path-set-1 then
RP/0/RSP0/CPU0:router(config-rpl-if)# drop
RP/0/RSP0/CPU0:router(config-rpl-if)# endif
RP/0/RSP0/CPU0:router(config-rpl)#
```

The contents of the **then** clause may be an arbitrary sequence of action statements.

The following example shows an **if** statement with two action statements:

```
RP/0/RSP0/CPU0:router(config-rpl)# if origin is igp then
RP/0/RSP0/CPU0:router(config-rpl-if)# set med 42
RP/0/RSP0/CPU0:router(config-rpl-if)# prepend as-path 73 5
RP/0/RSP0/CPU0:router(config-rpl-if)# endif
RP/0/RSP0/CPU0:router(config-rpl)#
```

The **if** command also permits an **else** clause to be executed if the expression is false, as follows:

```
RP/0/RSP0/CPU0:router(config-rpl)# if med eq 200 then
RP/0/RSP0/CPU0:router(config-rpl-if)# set community (12:34) additive
RP/0/RSP0/CPU0:router(config-rpl-if)# else
RP/0/RSP0/CPU0:router(config-rpl-else)# set community (12:56) additive
RP/0/RSP0/CPU0:router(config-rpl-else)# endif
RP/0/RSP0/CPU0:router(config-rpl)#
```

The routing policy language (RPL) also provides syntax using the **elseif** command to string together a sequence of tests, as shown in the following example:

```
RP/0/RSP0/CPU0:router(config-rpl)# if med eq 150 then
RP/0/RSP0/CPU0:router(config-rpl-if)# set local-preference 10
RP/0/RSP0/CPU0:router(config-rpl-if)# elseif med eq 200 then
RP/0/RSP0/CPU0:router(config-rpl-elseif)# set local-preference 60
RP/0/RSP0/CPU0:router(config-rpl-elseif)# elseif med eq 250 then
RP/0/RSP0/CPU0:router(config-rpl-elseif)# set local-preference 110
RP/0/RSP0/CPU0:router(config-rpl-elseif)# else
RP/0/RSP0/CPU0:router(config-rpl-else)# set local-preference 0
RP/0/RSP0/CPU0:router(config-rpl-else)# endif
RP/0/RSP0/CPU0:router(config-rpl)#
```

The statements within an **if** statement may themselves be **if** statements, as shown in this example:

```
RP/0/RSP0/CPU0:router(config-rpl)# if community matches-any (12:34, 56:78) then
```

```
RP/0/RSP0/CPU0:router(config-rpl-if)# if med eq 150 then
RP/0/RSP0/CPU0:router(config-rpl-if)# drop
RP/0/RSP0/CPU0:router(config-rpl-if)# endif
RP/0/RSP0/CPU0:router(config-rpl-if)# set local-preference 100
RP/0/RSP0/CPU0:router(config-rpl-if)# endif
RP/0/RSP0/CPU0:router(config-rpl)#
```

The policy configuration shown sets the value of the local preference attribute to 100 on any route that has a community value of 12:34 or 56:78 associated with it. However, if any of these routes has a Multi Exit Descriminator (MED) value of 150, then each route with both the community value of 12:34 or 56:78 and a MED of 150 is dropped.

if route-aggregated

To match the aggregated routes from the other routes, use the **if route-aggregated** command in route policy configuration mode.

if route-aggregated

Syntax Description

route-aggregated	Checks if route is an aggregation of multiple routes.

Command Default

No default behavior or values

Command Modes

Route-policy configuration

Command History

Release	Modification
Release 5.2.0	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
route-policy	read, write

Examples

This example shows how to match the aggregated routes from other routes:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# route-policy  route-policy atomic_agg
RP/0/RSP0/CPU0:router(config-rpl)# if route-aggregated then
RP/0/RSP0/CPU0:router(config-rpl-if)# set extcommunity rt (1:1)
RP/0/RSP0/CPU0:router(config-rpl-if)# endif
RP/0/RSP0/CPU0:router(config-rpl)# end-policy
```

is-best-path

To tag the path selected as the best path use the**is-best-path** command in route policy configuration mode.

is-best-path

Syntax Description

is-best-path	Checks and tags the path selected as best-path.

Command Default

No default behavior or values.

Command Modes

Route-policy configuration

Command History

Release	Modification
Release 5.3.2	This command was introduced.

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes 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
route-policy	read, write

Example

```
  RP/0/RSP0/CPU0:router(config)# route-policy
  WORD  Route Policy name
  RP/0/RSP0/CPU0:router(config)# route-policy sample
  RP/0/RSP0/CPU0:router(config-rpl)# if destination i
  in   is-backup-path  is-best-external  is-best-path

  if destination is-best-path then
    set community community
  endif
end-policy
!
RP/0/RSP0/CPU0:router# sh version
Wed Jul  8 16:08:34.286 IST
Cisco IOS XR Software, Version 5.3.2.14I[EnXR]
Copyright (c) 2015 by Cisco Systems, Inc.
    Built on Fri Jun 26 17:35:45 IST 2015
    By router in RP/0/RSP0/CPU0
```

is-backup-path

To tag all the paths equal to the back up path use, **is-backup-path** command in route policy configuration mode.

is-backup-path

Syntax Description

is-backup-path	Checks and tags the path selected as backup path.

Command Default

No default behavior or values.

Command Modes

Route-policy configuration

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
route-policy	read, write

Example

```
RP/0/RSP0/CPU0:router(config)# route-policy
  WORD  Route Policy name
RP/0/RSP0/CPU0:router(config)# route-policy sample
RP/0/RSP0/CPU0:router(config-rpl)# if destination i
in      is-backup-path  is-best-external  is-best-path

RP/0/RSP0/CPU0:router(config)# route-policy
WORD  Route Policy name
RP/0/RSP0/CPU0:router(config)# route-policy sample
RP/0/RSP0/CPU0:router(config-rpl)# if destination i
in     is-backup-path  is-best-external  is-best-path
```

is-multi-path

To tag all the paths equal to the best path based on multi-path context use, **is-multi-path** command in route policy configuration mode.

is-multi-path

Syntax Description

is-multi-path	Checks and tag all the path equal to the as best-path.

Command Default

No default behavior or values.

Command Modes

Route-policy configuration

Command History

Release	Modification
Release 5.3.2	This command was introduced.

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes 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
route-policy	read, write

Example

```
RP/0/RSP0/CPU0:router(config)#route-policy
WORD  Route Policy name
RP/0/RSP0/CPU0:router(config)#route-policy sample
RP/0/RSP0/CPU0:router(config-rpl)#if destination i
in             is-backup-path  is-best-external  is-best-path

is-multi-path
RP/0/RSP0/CPU0:router(config-rpl)#if destination is-
is-backup-path  is-best-external  is-best-path  is-multi-path
RP/0/RSP0/CPU0:router(config-rpl)#if destination is-best-path then
RP/0/RSP0/CPU0:router(config-rpl-if)#set l
label             label-index  label-mode  level
community  lsm-root
RP/0/RSP0/CPU0:router(config-rpl-if)#set community community
RP/0/RSP0/CPU0:router(config-rpl-if)#endif
RP/0/RSP0/CPU0:router(config-rpl)#end-policy
RP/0/RSP0/CPU0:router(config)#commit
Wed Jul  8 16:08:23.436 IST
```

local-preference

To compare the local-preference attribute of a BGP route to an integer value or a parameterized value, use the local-preference command in route-policy configuration mode.

local-preference {**eq** | **is** | **ge** | **le**} {*numberparameter*}

Syntax Description

eq | **is** | **ge** | **le**	Equal to; exact match; greater than or equal to; less than or equal to.
number	Value assigned to a 32-bit unsigned integer. Range is 0 to 4294967295.
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 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **local- preference** command as a conditional expression within an **if** statement to compare the local-preference attribute to an integer value or a parameterized value.

Note For a list of all conditional expressions available within an **if** statement, see the **if** command.

The MED is a 32-bit unsigned integer. The **eq** operation compares the local-preference to either a static value or a parameterized value passed to a parameterized policy for equality with that value. A greater than or equal to comparison can also be done with the **ge** operator, and a less than or equal to comparison can be performed using the **le** operator.

Examples

The following example shows that if the **local-preference** is 10, local-preference is set to 100:

```
RP/0/RSP0RP0/CPU0:router(config-rpl)# if local-preference eq 10 then
RP/0/RSP0RP0/CPU0:router(config-rpl-if)# set weight 100
RP/0/RSP0RP0/CPU0:router(config-rpl-if)# endif
RP/0/RSP0RP0/CPU0:router(config-rpl)#
```

large-community is-empty

To check if a route has no large community attributes associated with it, use the **large-community is-empty** command in route-policy configuration mode.

large-community is-empty

Syntax Description

This command has no arguments or keywords.

Command Default

No default behavior or values

Command Modes

Route-policy configuration

Task ID

Task ID	Operations
route-policy	read, write

Command History

Release	Modification
Release 6.3.1	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.

Use the **large community is-empty** command as a conditional expression within an **if** statement to check if a route has community attributes associated with it.

Note For a list of all conditional expressions available within an **if** statement, see the **if** command.

This command takes no arguments and evaluates to true only if the route has no community attributes associated with it.

Task ID

Task ID	Operations
route-policy	read, write

The following example shows using the **large-community is-empty** clause to filter routes that do not have the large-community attribute set.

```
RP/0/RSP0/CPU0:router#config
RP/0/RSP0/CPU0:router(config)#route-policy lrg_comm_rp4
RP/0/RSP0/CPU0:router(config-rpl)#if large-community is-empty then
RP/0/RSP0/CPU0:router(config-rpl)#set local-preference 104
```

```
RP/0/RSP0/CPU0:router(config-rpl)#endif
RP/0/RSP0/CPU0:router(config-rpl)#end-policy
```

large-community matches-any

To configure the route policy to match any elements of a large-community set, use the **large-community matches-any** command in route-policy configuration mode.

`large-community matches-any` { *large-community-set-name or inline-large-community-set* | `parameter` }

Syntax Description

large-community-set-name	Name of a large community set.
inline-large-community-set	Inline large community set. The inline large community set must be enclosed in parentheses.
parameter	Parameter name. The parameter name must be preceded with a "$."

Command Default

No default behavior or values

Command Modes

Route-policy configuration

Task ID

Task ID	Operations
route-policy	read, write

Command History

Release	Modification
Release 6.3.1	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.

Use the **large-community matches-any** command as a conditional expression within an **if** statement in the route policy statements to match any element of a large-community set.

Note For a list of all conditional expressions available within an **if** statement, see the **if** command.

The large communities 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.

In route-policy statements, each integer in the BGP large community can be replaced by any of the following expressions:

- [x..y] — This expression specifies a range between x and y, inclusive.
- * — This expression stands for any number.

- peeras — This expression is replaced by the AS number of the neigbhor from which the community is received or to which the community is sent, as appropriate.
- not-peeras — This expression matches any number other than the peeras.
- private-as — This expression specifies any number in the private ASN range: [64512..65534] and [4200000000..4294967294].

Note The peeras and not-peeras expressions can only be used in large-community match statements that appear in route policies that are applied at the neighbor-in or neighbor-out attach points.

IOS regular expression (ios-regex) and DFA style regular expression (dfa-regex) can be used in any of the large-community policy match statements. For example, the IOS regular expression ios-regex '^5:.*:7$' is equivalent to the expression 5:*:7.

Task ID

Task ID	Operations
route-policy	read, write

Examples

The following example shows how to configure a route policy to match any element of a large -community set. This is a boolean condition and returns true if any of the large communities in the route match any of the large communities in the match condition.

```
RP/0/RSP0/CPU0:router#config
RP/0/RSP0/CPU0:router(config)#route-policy elbonia
RP/0/RSP0/CPU0:router(config-rpl)#if large-community matches-any (1:2:3, 4:5:*) then
RP/0/RSP0/CPU0:router(config-rpl)#set local-preference 94
RP/0/RSP0/CPU0:router(config-rpl)#endif
RP/0/RSP0/CPU0:router(config-rpl)#end-policy
```

large-community matches-every

To configure the route policy to match every element of a large-community set, use the **large-community matches-every** command in route-policy configuration mode.

```
large-community matches-every { large-community-set-name or inline-large-community-set
|parameter }
```

Syntax Description

large-community-set-name	Name of a large community set.
inline-large-community-set	Inline large community set. The inline large community set must be enclosed in parentheses.
parameter	Parameter name. The parameter name must be preceded with a "$."

Command Default

No default behavior or values

Command Modes

Route-policy configuration

Task ID

Task ID	Operations
route-policy	read, write

Command History

Release	Modification
Release 6.3.1	This command was introduced.

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.

Use the **large-community matches-every** command as a conditional expression within an **if** statement in the route policy statements to match every element of a large-community set.

Note For a list of all conditional expressions available within an **if** statement, see the **if** command.

The large communities 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.

In route-policy statements, each integer in the BGP large community can be replaced by any of the following expressions:

- [x..y] — This expression specifies a range between x and y, inclusive.
- * — This expression stands for any number.
- peeras — This expression is replaced by the AS number of the neigbhor from which the community is received or to which the community is sent, as appropriate.

- not-peeras — This expression matches any number other than the peeras.
- private-as — This expression specifies any number in the private ASN range: [64512..65534] and [4200000000..4294967294].

Note The peeras and not-peeras expressions can only be used in large-community match statements that appear in route policies that are applied at the neighbor-in or neighbor-out attach points.

IOS regular expression (ios-regex) and DFA style regular expression (dfa-regex) can be used in any of the large-community policy match statements. For example, the IOS regular expression ios-regex '^5:.*:7$' is equivalent to the expression 5:*:7.

Task ID

Task ID	Operations
route-policy	read, write

The following example shows how to configure a route policy where every match specification in the statement must be matched by at least one large community in the route.

```
RP/0/RSP0/CPU0:router#config
RP/0/RSP0/CPU0:router(config)#route-policy bob
RP/0/RSP0/CPU0:router(config-rpl)#if large-community matches-any (*:*:3, 4:5:*) then
RP/0/RSP0/CPU0:router(config-rpl)#set local-preference 94
RP/0/RSP0/CPU0:router(config-rpl)#endif
RP/0/RSP0/CPU0:router(config-rpl)#end-policy
```

In this example, routes with these sets of large communities return TRUE:

- (1:1:3, 4:5:10)
- (4:5:3) —This single large community matches both specifications.
- (1:1:3, 4:5:10, 7:6:5)

Routes with the following set of large communities return FALSE:

(1:1:3, 5:5:10)—The specification (4:5:*) is not matched.

large-community matches-within

To configure a route policy to match within a large community set, use the **large-community matches-within** command in route-policy configuration mode.

```
large-community matches-within { large-community-set-name or inline-large-community-set | parameter }
```

Syntax Description

large-community-set-name	Name of a large community set.
inline-large-community-set	Inline large community set. The inline large community set must be enclosed in parentheses.
parameter	Parameter name. The parameter name must be preceded with a "$."

Command Default

No default behavior or values

Command Modes

Route-policy configuration

Task ID

Task ID	Operations
route-policy	read, write

Command History

Release	Modification
Release 6.3.1	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.

This command is similar to the **large-community matches-any** command but every large community in the route must match at least one match specification. Note that if the route has no large communities, then it matches.

When large communities are specified in other commands, 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.

In route-policy statements, each integer in the BGP large community can be replaced by any of the following expressions:

- [x..y] — This expression specifies a range between x and y, inclusive.
- * — This expression stands for any number.
- peeras — This expression is replaced by the AS number of the neigbhor from which the community is received or to which the community is sent, as appropriate.
- not-peeras — This expression matches any number other than the peeras.

- private-as — This expression specifies any number in the private ASN range: [64512..65534] and [4200000000..4294967294].



Note The peeras and not-peeras expressions can only be used in large-community match statements that appear in route policies that are applied at the neighbor-in or neighbor-out attach points.

IOS regular expression (ios-regex) and DFA style regular expression (dfa-regex) can be used in any of the large-community policy match statements. For example, the IOS regular expression ios-regex '^5:.*:7$' is equivalent to the expression 5:*:7.

Task ID

Task ID	Operations
route-policy	read, write

The following example shows how to configure a route policy to match within a large community set.

```
RP/0/RSP0/CPU0:router#config
RP/0/RSP0/CPU0:router(config)#route-policy bob
RP/0/RSP0/CPU0:router(config-rpl)#if large-community matches-within (*:*:3, 4:5:*) then
RP/0/RSP0/CPU0:router(config-rpl)#set local-preference 103
RP/0/RSP0/CPU0:router(config-rpl)#endif
RP/0/RSP0/CPU0:router(config-rpl)#end-policy
```

In this example, routes with these sets of large communities return TRUE:

- (1:1:3, 4:5:10)
- (4:5:3)
- (1:2:3, 6:6:3, 9:4:3)

Routes with the following set of large communities return FALSE:

(1:1:3, 4:5:10, 7:6:5) —The large community (7:6:5) does not match

large-community-set

To define a set of large-communities, use the **large-community-set** command in global configuration mode. To remove the large-community set, use the **no** form of this command.

large-community-set *name*
no large-community-set *name*

Syntax Description

name	Name of the large-community set. Named large-community sets are used in route-policy match and set statements.

Command Default

No default behavior or values

Command Modes

global configuration

Command History

Release	Modification
Release 6.3.1	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.

The large communities 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.

In route-policy statements, each integer in the BGP large community can be replaced by any of the following expressions:

- [x..y] — This expression specifies a range between x and y, inclusive.
- * — This expression stands for any number.
- private-as — This expression specifies any number in the private ASN range: [64512..65534] and [4200000000..4294967294].

Task ID

Task ID	Operations
route-policy	read, write

Examples

This example shows how to create a named large-community set:

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#large-community-set catbert
RP/0/RSP0/CPU0:router(config-largecomm)#1:2:3,
RP/0/RSP0/CPU0:router(config-largecomm)#[5..9]:2:3
RP/0/RSP0/CPU0:router(config-largecomm)#1:3:*
RP/0/RSP0/CPU0:router(config-largecomm)#end-set
```

med

To compare the Multi Exit Discriminator (MED) to an integer value or a parameterized value or compare the MED attribute of a BGP route to an integer value, use the **med** command in route-policy configuration mode.

med {**eq** | **is** | **ge** | **le**} {*numberparameter*}

Syntax Description

eq | **is** | **ge** | **le**	Equal to; exact match; greater than or equal to; less than or equal to.
number	Value assigned to a 32-bit unsigned integer. Range is 0 to 4294967295.
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 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **med** command as a conditional expression within an **if** statement to compare the MED to an integer value or a parameterized value.

Note For a list of all conditional expressions available within an **if** statement, see the **if** command.

The MED is a 32-bit unsigned integer. The **eq** operation compares the MED to either a static value or a parameterized value passed to a parameterized policy for equality with that value. A greater than or equal to comparison can also be done with the **ge** operator, and a less than or equal to comparison can be performed using the **le** operator.

Task ID

Task ID	Operations
route-policy	read, write

Examples

The following example shows that if the **med** commands match, the local preference is set to 100:

```
RP/0/RSP0/CPU0:router(config-rpl)# if med eq 10 then
RP/0/RSP0/CPU0:router(config-rpl-if)# set local-preference 100
RP/0/RSP0/CPU0:router(config-rpl-if)# endif
```

```
RP/0/RSP0/CPU0:router(config-rpl)#
```

next-hop in

To compare the next-hop associated with the route to data contained in either an inline or a named prefix set, use the **next-hop in** command in route-policy configuration mode.

next-hop in {*prefix-set-nameinline-prefix-setparameter*}

Syntax Description

prefix-set-name	Name of a prefix set.
inline-prefix-set	Inline prefix set. The inline prefix set must be enclosed in parentheses.
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 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **next-hop in** command as a conditional expression within an **if** statement to compare the next-hop associated with the route to data contained in either an inline or a named prefix set. The result is true if any value in the prefix set matches the next-hop of the route. A comparison that refers to a named prefix set that has no elements in it returns false.

Note For a list of all conditional expressions available within an **if** statement, see the **if** command.

The next-hop is an IPv4 address entered as a dotted-decimal or an IPv6 address entered as a colon-separated hexadecimal.

Task ID

Task ID	Operations
route-policy	read, write

Examples

The following example shows that if the **next-hop in** commands match, the local preference is set to 100

```
RP/0/RSP0/CPU0:router(config-rpl)# if next-hop in some-prefix-set then
RP/0/RSP0/CPU0:router(config-rpl-if)# if next-hop in (10.0.0.5, fe80::230/64) then
RP/0/RSP0/CPU0:router(config-rpl-if)# set local-preference 0
```

```
RP/0/RSP0/CPU0:router(config-rpl-if)# endif
RP/0/RSP0/CPU0:router(config-rpl)#
```

orf prefix in

To configure an outbound route filter (ORF), use the **orf prefix in** command in route-policy configuration mode.

orf prefix in {*prefix-set-nameinline-prefix-set*}

Syntax Description

prefix-set-name	Name of a prefix set.
inline-prefix-set	Inline prefix set. The inline prefix set must be enclosed in parentheses.

Command Default

No default behavior or values

Command Modes

Route-policy 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **orf prefix in** command to match a prefix in a prefix set or an inline prefix set.

This command takes either a named prefix set or an inline prefix set value as an argument. It returns true if the destination NLRI matches any entry in the prefix set. An attempt to match destination using a prefix set that is defined but contains no elements returns false.

This command is used in the context of the orf route-policy attach point in BGP. The destination of a route is also known in Border Gateway Protocol (BGP) as its network-layer reachability information (NLRI). It comprises a prefix value and a mask length. The routing policy language (RPL) provides one operation on prefixes, testing them for matching against a list of prefix-match specifications using the **in** operator.

Examples

In the following example, the prefix set orfpreset1 and the route policy named orfpolicy are defined. Next, the orfpolicy is applied to the neighbor orf attach point.

If the prefix of the route matches any of the prefixes specified in orfpreset1 (211.105.1.0/24, 211.105.5.0/24, 211.105.11.0/24), then the prefix is dropped. If the prefix matches in(211.105.3.0/24, 211.105.7.0/24, 211.105.13.0/24), then the prefix is accepted. In addition to this inbound filtering, BGP sends these prefix entries to the upstream neighbor indicating a permit or deny so that the neighbor can make the same filter updates.

```
RP/0/RSP0/CPU0:router(config)# prefix-set orfpreset1
RP/0/RSP0/CPU0:router(config-pfx)# 211.105.1.0/24,
RP/0/RSP0/CPU0:router(config-pfx)# 211.105.5.0/24,
RP/0/RSP0/CPU0:router(config-pfx)# 211.105.11.0/24
RP/0/RSP0/CPU0:router(config-pfx)# end-set
!
!
RP/0/RSP0/CPU0:router(config)# route-policy orfpolicy
```

```
RP/0/RSP0/CPU0:router(config-rpl)# if orf prefix in orfpreset1 then
RP/0/RSP0/CPU0:router(config-rpl-if)# drop
RP/0/RSP0/CPU0:router(config-rpl-if)# endif
RP/0/RSP0/CPU0:router(config-rpl)# if orf prefix in (211.105.3.0/24, 211.105.7.0/24,
211.105.13.0/24) then
RP/0/RSP0/CPU0:router(config-rpl-if)# pass
RP/0/RSP0/CPU0:router(config-rpl-if)# endif
RP/0/RSP0/CPU0:router(config-rpl)# end-policy
!
!
RP/0/RSP0/CPU0:router(config)# router bgp 2
RP/0/RSP0/CPU0:router(config-bgp)# neighbor 1.1.1.1
RP/0/RSP0/CPU0:router(config-bgp-nbr)# remote-as 3
RP/0/RSP0/CPU0:router(config-bgp-nbr)# address-family ipv4 unicast
RP/0/RSP0/CPU0:router(config-bgp-nbr-af)# orf route-policy orfpolicy
```

Related Commands

Command	Description
orf	Specifies BGP ORF and inbound filtering criteria.

origin is

To match a specific origin type, use the **origin is** command in route-policy configuration mode.

origin is {**igp** | **egp** | **incomplete***parameter*}

Syntax Description

igp	Specifies Interior Gateway Protocol.
egp	Specifies Exterior Gateway Protocol.
incomplete	Specifies that Border Gateway Protocol (BGP) first learned the route by means other than BGP or Interior Gateway Protocol (IGP); for example, the route is learned through configuration.
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 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **origin is** command as a conditional expression within an **if** statement to test the value of the origin attribute.

Note For a list of all conditional expressions available within an **if** statement, see the **if** command.

The origin of a BGP route is an enumeration; it is **igp** , **egp** , or **incomplete** .

This command can be parameterized.

Task ID

Task ID	Operations
route-policy	read, write

Examples

In the following example, the origin is tested within an **if** statement to learn if it is either **igp** or **egp** :

```
RP/0/RSP0/CPU0:router(config-rpl)# if origin is igp or origin is egp then
```

In the following example, a parameter is used to match a specific origin type:

```
RP/0/RSP0/CPU0:router(config)# route-policy bar($origin)
RP/0/RSP0/CPU0:router(config-rpl)# if origin is $origin then
RP/0/RSP0/CPU0:router(config-rpl-if)# set med 20
RP/0/RSP0/CPU0:router(config-rpl-if)# endif
RP/0/RSP0/CPU0:router(config-rpl)#
```

ospf-area

To match a specific ospf area, use the **ospf-area** command in route-policy configuration mode.

ospf-area [**all-paths**] {**in** | **is**}

Syntax Description

is	Specify the explicit *area-id* .
in	Specify a list of *area-id* or *area-set*. Multiple areas can be specified separated by a comma (,).
all-paths	Used for routes with multiple paths. A match is made if area for every path of the route is configured in the route-policy.

Command Default

None

Command Modes

Route-policy configuration

Command History

Release	Modification
Release 5.2.2	This command was introduced.

Usage Guidelines

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

The route policy define by using **ospf-area** is useful in redistributing routes from a specific area of a routing domain into OSPF. After the route policy is crated, use the **redistribute ospf route-policy** command for route redistribution.

Task ID

Task ID	Operations
route-policy	read, write

Example

In the following example, an explicit area is specified as the matching criteria.

```
RP/0/RSP0/CPU0:router(config-rpl)# if ospf-area is 10 then pass else drop endif
```

In the following example, a collection of areas is specified as the matching criteria.

```
RP/0/RSP0/CPU0:router(config-rpl)# if ospf-area in (5,6,255.255.10.2) then pass else drop
endif
```

In the following example, an area set is specified as the matching criteria. As a pre-requisite, the area set must be defined.

```
RP/0/RSP0/CPU0:router(config)# ospf-area-set S1
RP/0/RSP0/CPU0:router(config-ospf-area)# 1 , 2.2.2.2 end-set
RP/0/RSP0/CPU0:router(config)# route-policy P1
RP/0/RSP0/CPU0:router(config-rpl)# if ospf-area in S1 then pass else drop endif
```

ospf-area-set

Defines an OSPF area set to be used in routing policy statements.

ospf-area-set*name* {*<0-4294967295>*<*A.B.C.D*> | **abort** | **end-set** | **exit** | **show**}
noospf-area-set*name*

Syntax Description

name	Name of the OSPF area set.
<0-4294967295>	32-bit decimal number to identify the set.
<A.B.C.D>	IPv4 Address used to identify the set, or the IPv4 address of the ACL.
abort	Exits the OSPF area set configuration without committing.
end-set	Exits the set configuration mode. You can commit the set after this option.
show	Displays the partial RPL configuration.

Command Default

No default behavior or values

Command Modes

Global configuration

Command History

Release	Modification
Release 5.1.1	This command was introduced.

Task ID

Task ID	Operations
route-policy	read, write

Example

The example shows how you can configure OSPF area sets with wildcards in routing policy.

```
RP/0/RSP0/CPU0:router(config)# ospf-area-set ospf_area_set_demo1
RP/0/RSP0/CPU0:router(config-ospf-area)# 10.0.0.1,
RP/0/RSP0/CPU0:router(config-ospf-area)# 3553
RP/0/RSP0/CPU0:router(config-ospf-area)# end-set

RP/0/RSP0/CPU0:router(config)# ospf-area-set ospf_area_set_demo2
RP/0/RSP0/CPU0:router(config-ospf-area)# 20.0.0.2,
RP/0/RSP0/CPU0:router(config-ospf-area)# 3673
RP/0/RSP0/CPU0:router(config-ospf-area)# end-set

RP/0/RSP0/CPU0:router(config)# route-policy use_ospf_area_set
RP/0/RSP0/CPU0:router(config-rpl)# if ospf-area in ospf-area-set* then set ospf-metric 200
RP/0/RSP0/CPU0:router(config-rpl-if)# elseif ospf-area in( 10.0.0.1, 10.0.0.2 )then set
```

```
ospf-metric 300
RP/0/RSP0/CPU0:router(config-rpl-elseif)# endif
RP/0/RSP0/CPU0:router(config-rpl)# end-policy
```

pass

To pass a route for further processing, use the **pass** command in route-policy configuration mode.

pass

Syntax Description

This command has no arguments or keywords.

Command Default

No default behavior or values

Command Modes

Route-policy 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **pass** command to signify that even though this route has not been modified, the user wants to continue executing in this policy block.

Note The **pass** command can be used as an action statement within an **if** statement. For a list of all action statements available within an **if** statement, see the **if** command.

When a policy block has finished executing, any route that has been modified in this policy block or has received a pass disposition in this policy block passes the policy and execution finishes for that policy. If this policy block is applied from within another policy block and the route is either passed or modified, then execution continues in the policy block that applied this policy block.

Task ID

Task ID	Operations
route-policy	read, write

Examples

The following example shows how to accept the route unconditionally without modifying it:

```
RP/0/RSP0/CPU0:router(config-rpl)# pass
```

This example accepts the route unconditionally, without modifying it, if the destination is in prefix-set permitted:

```
RP/0/RSP0/CPU0:router(config-rpl)# if destination in permitted then
RP/0/RSP0/CPU0:router(config-rpl-if)# pass
RP/0/RSP0/CPU0:router(config-rpl-if)# endif
```

```
RP/0/RSP0/CPU0:router(config-rpl)#
```

path-type is

To match path types, use the **path-type is** command in route-policy configuration mode.

path-type is {**ibgp** | **ebgp***parameter*}

Syntax Description

ibgp	Specifies an internal BGP path.
ebgp	Specifies an external BGP path.
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 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **path-type is** command as a conditional expression within an **if** statement to match path types.

Note For a list of all conditional expressions available within an **if** statement, see the **if** command.

Task ID

Task ID	Operations
route-policy	read, write

Examples

In the following example, if the path is an external BGP path the route is accepted:

```
RP/0/RSP0/CPU0:router(config)# route-policy policy_A
RP/0/RSP0/CPU0:router(config-rpl)# if path-type is ebgp then
RP/0/RSP0/CPU0:router(config-rpl-if)# pass
RP/0/RSP0/CPU0:router(config-rpl-if)# else
RP/0/RSP0/CPU0:router(config-rpl-else)# drop
RP/0/RSP0/CPU0:router(config-rpl-if)# endif
RP/0/RSP0/CPU0:router(config-rpl)# end-policy
```

policy-global

To define global parameters and enter global parameter configuration mode, use the **policy-global** command in global configuration mode. To remove global parameters, use the **no** form of this command.

policy-global
no policy-global

Syntax Description

This command has no arguments or keywords.

Command Default

No default behavior or values

Command Modes

Global 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **policy-global** command to define global parameters and enter global parameter configuration mode.

RPL supports the definition of systemwide global parameters that can be used inside a policy definition. The global parameter values can be used directly inside a policy definition similar to the local parameters of parameterized policy. When a parameterized policy has a parameter name "collision" with a global parameter name, parameters local to policy definition take precedence, effectively 'masking off' global parameters. In addition, a validation mechanism is in place to prevent the deletion of a particular global parameter if it is referred by any policy. For more information on global parameters and parameterization, see the *Implementing Routing Policy* module of the *Routing Configuration Guide for Cisco ASR 9000 Series Routers*

Task ID

Task ID	Operations
route-policy	read, write

Examples

The following example shows how to configure global parameters:

```
RP/0/RSP0/CPU0:router(config)# policy-global
RP/0/RSP0/CPU0:router(config-rp-gl)# glbpathtype 'ebgp'
RP/0/RSP0/CPU0:router(config-rp-gl)# glbtag '100'
RP/0/RSP0/CPU0:router(config-rp-gl)# end-global
```

In the following example, the *globalparam* argument makes use of the global parameters gbpathtype and glbtag defined above and is defined for a nonparameterized policy:

```
RP/0/RSP0/CPU0:router(config)# route-policy globalparam
RP/0/RSP0/CPU0:router(config-rpl)# if path-type is $glbpathtype then
```

```
RP/0/RSP0/CPU0:router(config-rpl)# set tag $glbtag
RP/0/RSP0/CPU0:router(config-rpl)# endif
RP/0/RSP0/CPU0:router(config-rpl)# end-policy
```

Related Commands

Command	Description
end-global, on page 1649	Ends the definition of global parameters.

prefix-set

To enter prefix set configuration mode and define a prefix set for contiguous and non-contiguous set of bits, use the **prefix-set** command in global configuration mode. To remove a named prefix set, use the **no** form of this command.

prefix-set *name*
no prefix-set *name*

Syntax Description

name	Name of a prefix set.

Command Default

None

Command Modes

Global 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **prefix-set** command to enter prefix set configuration mode and define a prefix set.

A prefix set is a comma-separated list of prefix match specifications. It holds IPv4 or IPv6 prefix match specifications, each of which has four parts: an address, a mask length, a minimum matching length, and a maximum matching length. The address is required, but the other three parts are optional. The address is a standard four-part, dotted-decimal numeric IPv4 address or a colon-separated hexadecimal IPv6 address. The mask length, if present, is a nonnegative decimal integer in the range from 0 to 32 for IPv4 prefixes or 0 to 128 for IPv6 prefixes following the address and separated from it by a slash. The optional minimum matching length follows the address and optional mask length and is expressed as the keyword **ge** (mnemonic for **g**reater than or **e**qual to), followed by a nonnegative decimal integer in the range from 0 to 32 for IPv4 or 0 to 128 for IPv6. The optional maximum matching length follows the rest and is expressed by the keyword **le** (mnemonic for **l**ess than or **e**qual to), followed by yet another nonnegative decimal integer in the range from 0 to 32 for IPv4 or 0 to 128 for IPv6. A syntactic shortcut for specifying an exact length for prefixes to match is the **eq** keyword, mnemonic for **eq**ual to.

If a prefix match specification has no mask length, then the default mask length is 32 for IPv4 or 128 for IPv6. The default minimum matching length is the mask length. If a minimum matching length is specified, then the default maximum matching length must be less than 32 for IPv4 prefixes or 128 for IPv6 prefixes. Otherwise, if neither a minimum nor maximum length is specified, the default maximum length is the mask length.

Task ID

Task ID	Operations
route-policy	read, write

Examples

The following example shows a prefix set named legal-ipv4-prefix-examples:

```
RP/0/RSP0/CPU0:router(config)# prefix-set legal-ipv4-prefix-examples
RP/0/RSP0/CPU0:router(config-pfx)# 10.0.1.1,
RP/0/RSP0/CPU0:router(config-pfx)# 10.0.2.0/24,
RP/0/RSP0/CPU0:router(config-pfx)# 10.0.3.0/24 ge 28,
RP/0/RSP0/CPU0:router(config-pfx)# 10.0.4.0/24 le 28,
RP/0/RSP0/CPU0:router(config-pfx)# 10.0.5.0/24 ge 26 le 30,
RP/0/RSP0/CPU0:router(config-pfx)# 10.0.6.0/24 eq 28
RP/0/RSP0/CPU0:router(config-pfx)# end-set
```

The first element of the prefix set matches only one possible value, 10.0.1.1/32 or the host address 10.0.1.1. The second element matches only one possible value, 10.0.2.0/24. The third element matches a range of prefix values, from 10.0.3.0/28 to 10.0.3.255/32. The fourth element matches a range of values, from 10.0.4.0/24 to 10.0.4.240/28. The fifth element matches prefixes in the range from 10.0.5.0/26 to 10.0.5.252/30. The sixth element matches any prefix of length 28 in the range from 10.0.6.0/28 through 10.0.6.240/28.

The following prefix set consists entirely of invalid prefix match specifications:

```
RP/0/RSP0/CPU0:router(config)# prefix-set INVALID-PREFIX-EXAMPLES
RP/0/RSP0/CPU0:router(config-pfx)# 10.1.1.1 ge 16,
RP/0/RSP0/CPU0:router(config-pfx)# 10.1.2.1 le 16,
RP/0/RSP0/CPU0:router(config-pfx)# 10.1.3.0/24 le 23,
RP/0/RSP0/CPU0:router(config-pfx)# 10.1.4.0/24 ge 33,
RP/0/RSP0/CPU0:router(config-pfx)# 10.1.5.0/25 ge 29 le 28
RP/0/RSP0/CPU0:router(config-pfx)# end-set
```

Neither the minimum length nor the maximum length is legal without a mask length. The maximum length must be at least the mask length. The minimum length must be less than 32, the maximum length of an IPv4 prefix. The maximum length must be equal to or greater than the minimum length.

The following example shows a valid IPv6 prefix set named legal-ipv6-prefix-examples:

```
RP/0/RSP0/CPU0:router(config)# prefix-set legal-ipv6-prefix-examples
RP/0/RSP0/CPU0:router(config-pfx)# 2001:0:0:1::/64,
RP/0/RSP0/CPU0:router(config-pfx)# 2001:0:0:2::/64,
RP/0/RSP0/CPU0:router(config-pfx)# 2001:0:0:3::/64,
RP/0/RSP0/CPU0:router(config-pfx)# 2001:0:0:4::/64
RP/0/RSP0/CPU0:router(config-pfx)# end-set
```

prepend as-path

To prepend the AS path with additional autonomous system numbers, use the **prepend as-path** command in route-policy configuration mode.

prepend as-path {*as-numberparameter* | **most-recent**} [*numberparameter*]

Syntax Description

as-number	Autonomous system number to prepend to the path. • 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.
parameter	Parameter name. The parameter name must be preceded with a "$."
most-recent	Specifies that the most recent autonomous system number should be prepended.
number	(Optional) Number of times the autonomous system number should be prepended. Range is 1 to 63.

Command Default

The default *number* is 1.

Command Modes

Route-policy configuration

Command History

Release	Modification
Release 3.7.2	This command was introduced.
Release 3.9.0	Asplain format for 4-byte Autonomous system numbers notation was supported.

Usage Guidelines

Use the **prepend as-path** command to prepend the AS path with additional autonomous system numbers.

Note The **prepend as-path** command can be used as an action statement within an **if** statement. For a list of all action statements available within an **if** statement, see the **if** command.

This command can take one or two arguments. The first argument (either a number or parameter) is the autonomous system number to prepend to the path. The optional second argument (either a number or parameter) is the number of times the autonomous system number should be prepended.

Task ID

Task ID	Operations
route-policy	read, write

Examples

The following example shows how to prepend the autonomous system number 666.1 to the AS path three times:

```
RP/0/RSP0/CPU0:router(config-rpl)# prepend as-path 666.1 3
```

The following example shows how to prepend the autonomous system number 666.0 to the AS path one time:

```
RP/0/RSP0/CPU0:router(config-rpl)# prepend as-path 666.0 1
```

protocol

To check the protocol that installs the route, use the **protocol** command in route-policy configuration mode.

protocol {**in**(*protocol-set)* | **is***protocol-name*}

Syntax Description

in (*protocol-set*)	Specifies a member of a set. The *protocol-set* argument accepts the following keywords within parentheses: • **bgp** —Border Gateway Protocol (BGP) • **connected** —Connected routes • **eigrp** —Enhanced Interior Gateway Routing Protocol (EIGRP) • **isis** —ISO Intermediate System-to-Intermediate System (IS-IS) • **ospf** —Open Shortest Path First (OSPF) • **ospfv3** —Open Shortest Path First version 3 (OSPFv3) • **rip** —Routing Information Protocol (RIP) • **static** —Static routes Keywords must be separated by a comma.
is *protocol-name*	Specifies a single protocol name, and accepted keywords are similar to the *protocol-set* argument.

Command Default

No default behavior or values

Command Modes

Route-policy 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **protocol** command as a conditional expression within an if statement to specify a protocol to install a route.

Use the **in** keyword to determine if a protocol listed in the *protocol-set* is the originator of the route being filtered.

Use the **is** keyword to determine if *protocol-name* is an exact match.

Note For a list of all conditional expressions available within an **if** statement, see the **if** command.

Task ID

Task ID	Operations
route-policy	read, write

Examples

The following example shows how to use the **protocol** command as a conditional expression within if statements:

```
RP/0/RSP0/CPU0:router(config)# route-policy rip1
RP/0/RSP0/CPU0:router(config-rpl)# if protocol in (connected, static) then
RP/0/RSP0/CPU0:router(config-rpl-if)# add rip-metric 2
RP/0/RSP0/CPU0:router(config-rpl-if)# elseif protocol is bgp 1 then
RP/0/RSP0/CPU0:router(config-rpl-elseif)# add rip-metric 3
RP/0/RSP0/CPU0:router(config-rpl-elseif)# elseif protocol is ospf 2 then
RP/0/RSP0/CPU0:router(config-rpl-elseif)# add rip-metric 4
RP/0/RSP0/CPU0:router(config-rpl-elseif)# else
RP/0/RSP0/CPU0:router(config-rpl-else)# add rip-metric 5
RP/0/RSP0/CPU0:router(config-rpl-else)# endif
RP/0/RSP0/CPU0:router(config-rpl)# end-policy

RP/0/RSP0/CPU0:router(config)# router rip
RP/0/RSP0/CPU0:router(config-rip)# interface GigabitEthernet0/1/0/1
RP/0/RSP0/CPU0:router(config-rip-if)# route-policy rip1 out
```

rd in

To compare the route distinguisher (RD) associated with the route to RDs contained in either a named or an inline RD set, use the **rd in** command in route-policy configuration mode.

rd in {*rd-set-nameinline-rd-setparameter*}

Syntax Description

rd-set-name	Name of an RD set.
inline-rd-set	Inline RD set. The inline RD set must be enclosed in parentheses.
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 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **rd in** command as a conditional expression within an **if** statement to match a destination entry in a named prefix set or inline prefix set.

Note For a list of all conditional expressions available within an **if** statement, see the **if** command.

This command takes either a named RD set or an inline RD set value as an argument. The condition returns true if the destination entry matches any entry in the RD set or inline RD set. An attempt to match an RD using an RD set that is defined but contains no elements returns false.

Task ID

Task ID	Operations
route-policy	read, write

Examples

The following example shows the **rd in** command with an inline RD set value as an argument:

```
RP/0/RSP0/CPU0:router(config)# route-policy
RP/0/RSP0/CPU0:router(config-rpl)# if rd in (128.1.0.0/16:100) then
RP/0/RSP0/CPU0:router(config-rpl-if)# pass
RP/0/RSP0/CPU0:router(config-rpl-if)# endif
RP/0/RSP0/CPU0:router(config-rpl)# end-policy
```

rd-set

To define a route distinguisher (RD) set and enter RD configuration mode, use the **rd-set** command in global configuration mode.

rd-set *name*
no rd-set *name*

Syntax Description

name	Name of an RD community set.

Command Default

No default behavior or values

Command Modes

Global 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **rd-set** command to create a set with RD elements and enter RD configuration mode. An RD set is a 64-bit value prepended to an IPv4 address to create a globally unique Border Gateway Protocol (BGP) VPN IPv4 address.

Note For *m*, the mask length is supported.

You can define RD values with the following commands:

- *a.b.c.d/m:**—BGP VPN RD in IPv4 format with a wildcard character. For example, 10.0.0.2/24.0:*.
- *a.b.c.d/m:n*—BGP VPN RD in IPv4 format with a mask. For example, 10.0.0.2/24:666.
- *a.b.c.d:** —BGP VPN RD in IPv4 format with a wildcard character. For example, 10.0.0.2:*.
- *a.b.c.d:n*—BGP VPN RD in IPv4 format. For example, 10.0.0.2:666.
- *asn:**—BGP VPN RD in ASN format with a wildcard character. For example, 10002:*.
- *asn:n*—BGP VPN RD in ASN format. For example, 10002:666.
- *x.y:**—BGP VPN RD in 4-byte ASN format with a wildcard character. For example, 10002.101:*.
- *x.y:n*—BGP VPN RD in 4-byte ASN format. For example, 10002.101:666.

Task ID

Task ID	Operations
route-policy	read, write

Examples

The following example shows how to create an RD set called my_rd_set:

```
RP/0/RSP0/CPU0:router(config)# rd-set my_rd_set
RP/0/RSP0/CPU0:router(config-rd)# 172.16.0.0/16:*,
RP/0/RSP0/CPU0:router(config-rd)# 172.17.0.0/16:100,
RP/0/RSP0/CPU0:router(config-rd)# 192:*,
RP/0/RSP0/CPU0:router(config-rd)# 192:100
RP/0/RSP0/CPU0:router(config-rd)# end-set
```

replace as-path

To replace a sequence of AS numbers or private AS numbers in the AS path with the configured local AS number, use the **replace as-path** command in route-policy configuration mode.

replace as-path {[*as-number-list parameter*] | **private-as**}

Syntax Description

as-number-list	(Optional) Sequence of AS numbers to replace. The sequence must be enclosed in single quotes (‘ ’). You can use 2-byte or 4-byte AS numbers. • The 2-byte value is entered as a 16-bit unsigned decimal value. The range is 0 to 65535. • The 4-byte value is entered as two 16-bit unsigned decimal values separated by a period. The range is 1.0 to 65535.65535.
parameter	(Optional) Parameter name. The parameter name must be preceded with a “$.”
private-as	Matches within the BGP private AS range. Range is from 64512 to 65534.

Command Default

None.

Command Modes

Route-policy 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **replace as-path** command to replace a sequence of AS numbers or private AS numbers in the AS path with the local AS numbers. For example, if the AS path is ‘67 65534 100 65533 5 78 89 90’ and the local AS number is 900, then:

```
replace as-path ‘5 78’
```

replaces‘5 78’ in the AS path with 900 (from the local AS), and the new path would be‘67 65534 100 65533 900 89 90’.

Consider following statement:

```
replace as-path private-as
```

Because 65534 and 65533 are within the private AS range, they are replaced with 900. The path is ‘67 900 100 900 5 78 89 90’. The length of the path remains the same.

The **replace as-path** command can be used as an action statement within an **if** statement. For a list of all action statements available within an **if** statement, see the **if** command.

Caution The **replace as-path** command changes the AS path content which can lead to routing loops.

Task ID

Task ID	Operations
route-policy	read, write

Examples

The following example shows how to use the **replace as-path** command to replace AS numbers in the AS path:

```
RP/0/RSP0/CPU0:router(config)# route-policy drop-as-1234
RP/0/RSP0/CPU0:router(config-rpl)# replace as-path '90 78 45 $asnum'
RP/0/RSP0/CPU0:router(config-rpl)# replace as-path private-as
RP/0/RSP0/CPU0:router(config-rpl)# replace as-path '9.9 7.89 14.15 $asnum'
RP/0/RSP0/CPU0:router(config-rpl)# replace as-path '9 89 14.15 $asnum'
```

remove as-path private-as

To remove BGP private AS numbers from as-path structure used by BGP, use the **remove as-path private-as** command under route policy configuration mode.

remove as-path private-as [entire-aspath]

Syntax Description

entire-aspath	(Optional) Removes the entire private autonomous system numbers from an autonomous system path only if all the autonomous systems in the path are private.

Command Default

No default behavior or values

Command Modes

Route-policy configuration

Command History

Release	Modification
Release 5.2.0	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
route-policy	read, write

Examples

This example shows how to remove BGP private AS numbers from as-path structure:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# route-policy rm_private_as
RP/0/RSP0/CPU0:router(config-rpl)# remove as-path private-as entire-aspath
RP/0/RSP0/CPU0:router(config-rpl)# end-policy
```

rib-has-route

To check if a route listed in the prefix set exists in the Routing Information Base (RIB), use the **rib-has-route** command in route-policy configuration mode.

rib-has-route { **async** | **in** } { *prefix-set-name inline-prefix-set parameter* }

Syntax Description

async	Enable the asynchronous (event-driven) route lookup
in	Specify the synchronous (polling-driven) implementation
prefix-set-name	Name of a prefix set.
inline-prefix-set	Inline prefix set. The inline prefix set must be enclosed in parentheses.
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 3.7.2	This command was introduced.
Release 7.4.1	The **async** keyword was added.

Usage Guidelines

The **rib-has-route-in** implementation is time-based, using a 60-second polling of the RIB to verify that at least one of the user-specified prefixes is present in the RIB. The **rib-has-route-async** implementation is an event-based way to track the prefixes in the RIB, so that if a prefix is added or removed from the RIB, BGP is notified and reruns the policy.

If routes are active, then they are advertised. Routes are considered active if they are already installed in the Routing Information Base (RIB).

The prefix sets used in the **rib-has-route** command contain two match specifications. The first is where an exact route match is requested (for example, 10.10.0.0/16 will match exactly one route) and the second is where a route match or any more-specific route match is allowed (for example, 10.10.0.0/16 ge 32 will match the 10.10.0.0/16 route and any longer prefix).

Use the **rib-has-route** command as a conditional expression within an **if** statement to check if there is an active route with a specific prefix contained in the RIB. If the statement reveals an active route that meets that criteria, additional actions are executed.

For a list of all conditional expressions available within an **if** statement, see the **if** command.

Task ID

Task ID	Operations
route-policy	read, write

Examples

In the following example, an **if** statement, with the **async** keyword, is used to track a prefix set 10.10.0.0/16 is in the RIB:

```
RP/0/RSP0/CPU0:router(config)# route-policy SAMPLE
RP/0/RSP0/CPU0:router(config-rpl)# if rib-has-route async (10.10.0.0/16 ge 16) then
RP/0/RSP0/CPU0:router(config-rpl-if)# pass
RP/0/RSP0/CPU0:router(config-rpl-if)# endif
RP/0/RSP0/CPU0:router(config-rpl)# end-policy
```

In the following example, an **if** statement is used to learn if a route contained in a prefix set 10.10.0.0/16 is in the RIB:

```
RP/0/RSP0/CPU0:router(config)# route-policy SAMPLE
RP/0/RSP0/CPU0:router(config-rpl)# if rib-has-route in (10.10.0.0/16 ge 16) then
RP/0/RSP0/CPU0:router(config-rpl-if)# pass
RP/0/RSP0/CPU0:router(config-rpl-if)# endif
RP/0/RSP0/CPU0:router(config-rpl)# end-policy
```

route-has-label

To check if there is a Multiprotocol Label Switching (MPLS) label in a route during redistribution, use the **route-has-label** command in route-policy configuration mode.

route-has-label

Syntax Description

This command has no arguments or keywords.

Command Default

No default behavior or values

Command Modes

Route-policy 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **route-has-label** command as a conditional expression within an **if** statement to check if there is an MPLS label in a route during redistribution.

For a list of all conditional expressions available within an **if** statement, see the **if** command.

Task ID

Task ID	Operations
route-policy	read, write

Examples

In the following example, an **if** statement learns if an MPLS label is present in a route:

```
RP/0/RSP0/CPU0:router(config-rpl)# if route-has-label then
RP/0/RSP0/CPU0:router(config-rpl-if)# pass
RP/0/RSP0/CPU0:router(config-rpl-if)# endif
RP/0/RSP0/CPU0:router(config-rpl)#
```

route-policy (RPL)

To define a route policy and enter route-policy configuration mode, use the **route-policy** command in global configuration mode. To remove a policy definition, use the **no** form of this command.

route-policy *name* [*(parameter1, parameter2, . . . , parameterN)*]
no route-policy *name*
(parameter1, parameter2, . . . , parameterN)

Syntax Description

name	Name of a route policy.
parameter	(Optional) Parameter name. The parameter name must be preceded with a "$." The *parameters* must be enclosed in parenthesis "()".

Command Default

No default behavior or values

Command Modes

Global 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **route-policy** command to define a route policy and enter route-policy configuration mode.

Policy definitions create named bundles of policy statements. A policy definition consists of the **route-policy** command followed by a name, a group of policy statements, and the **end-policy** command.

The policy name serves as a handle for binding the policy to protocols.

Task ID

Task ID	Operations
route-policy	read, write

Examples

The following example shows a simple policy named drop-everything that drops any route it encounters:

```
RP/0/RSP0/CPU0:router(config)# route-policy drop-everything
RP/0/RSP0/CPU0:router(config-rpl)# drop
RP/0/RSP0/CPU0:router(config-rpl)# end-policy
```

Policies may also refer to other policies such that common blocks of policy can be reused. This reference to other policies is accomplished by using the **apply** command. The following is a simple example:

```
RP/0/RSP0/CPU0:router(config)# route-policy drop-as-1234
RP/0/RSP0/CPU0:router(config-rpl)# if as-path passes-through '1234' then
RP/0/RSP0/CPU0:router(config-rpl-if)# apply check-communities
RP/0/RSP0/CPU0:router(config-rpl-if)# else
RP/0/RSP0/CPU0:router(config-rpl-else)# pass
RP/0/RSP0/CPU0:router(config-rpl-else)# endif
RP/0/RSP0/CPU0:router(config-rpl)# end-policy
```

The **apply** command indicates that the policy check-communities should be executed if the route under consideration passed through autonomous system 1234 before it was received. If so, the communities of the route are checked, and based on the findings the route may be accepted unmodified, accepted with changes, or dropped.

Related Commands

Command	Description
end-policy, on page 1650	Ends the definition of a route policy.

route-type is

To match route types when redistribution is being performed into Border Gateway Protocol (BGP), Open Shortest Path First (OSPF), or Integrated Intermediate System-to-Intermediate System (IS-IS), use the **route-type is** command in route-policy configuration mode.

route-type is {**local** | **interarea** | **internal** | **type-1** | **type-2** | **level-l** | **level-2***parameter*}

Syntax Description

local	Uses a local value to match locally generated BGP routes.
interarea	Uses an interarea value to match IS-IS interarea routes.
internal	Uses an internal value to match OSPF intra- and interarea routes.
type-1	Uses a Type 1 value to match Type 1 OSPF routes.
type-2	Uses a Type 2 value to match Type 2 OSPF routes.
level-1	Uses a Level 1 value to match Level 1 IS-IS routes.
level-2	Uses a Level 2 value to match Level 2 IS-IS routes.
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 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **route-type is** command as a conditional expression within an **if** statement to compare route types when redistribution is being performed into BGP, OSPF, or IS-IS.

Note For a list of all conditional expressions available within an **if** statement, see the **if** command.

The valid keywords are **local** , **internal** , **interarea** , **type-1** , **type-2** , **level-1** , and **level-2** . A parameterized value that fills in one of these values may also be used. The **local** value is used to match locally generated BGP routes. The internal value is used to match OSPF intra- and interarea routes. The **type-1** and **type-2** values are used to match Type 1 and Type 2 OSPF external routes. The **level-1** , **level-2** , and **interarea** values are used to match IS-IS routes of those respective types.

Because the route type is a matching operator, it appears in conditional clauses of **if** and **then** statements.

Task ID

Task ID	Operations
route-policy	read, write

Examples

In the following example, non-local routes are dropped:

```
RP/0/RSP0/CPU0:router(config)# route-policy policy_A
RP/0/RSP0/CPU0:router(config-rpl)# if route-type is local then
RP/0/RSP0/CPU0:router(config-rpl-if)# pass
RP/0/RSP0/CPU0:router(config-rpl-if)# else
RP/0/RSP0/CPU0:router(config-rpl-else)# drop
RP/0/RSP0/CPU0:router(config-rpl-if)# endif
RP/0/RSP0/CPU0:router(config-rpl)# end-policy
```

rpl editor

To set the default routing policy language (RPL) editor, use the **rpl editor** command in global configuration mode.

rpl editor {**nano** | **emacs** | **vim**}

Syntax Description

nano	Sets the default RPL editor to GNU nano.
emacs	Sets the default RPL editor to EMACS.
vim	Sets the default RPL editor to VIM.

Command Default

The Nano editor is the default.

Command Modes

Global configuration

Command History

Release	Modification
Release 3.7.2	This command was introduced.
Release 7.11.1	This command was deprecated.

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
route-policy	read, write

Examples

In the following example, the default RPL editor is set to Nano:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# rpl editor nano
```

In the following example, the default RPL editor is set to EMACS:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# rpl editor emacs
```

In the following example, the default RPL editor is set to VIM:

```
RP/0/RSP0/CPU0:router# configure
```

```
RP/0/RSP0/CPU0:router(config)# rpl editor vim
```

rpl maximum

To configure system limits on the routing policy subsystem, use the **rpl maximum** command in global configuration mode.

rpl maximum {**lines** | **policies**} *number*

Syntax Description

lines *number*	Configures the number of lines of configuration limit. Range is from 1 to 131072.
policies *number*	Configures the number of policies limit. Range is from 1 to 5000.

Command Default

lines *number* : 65536

policies *numbers* : 3500

Command Modes

global 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **rpl maximum** command to configure system limits on the routing policy subsystem. As such, **rpl maximum** configuration lines do not appear as statements within a routing policy. This command places resource limits on the routing policy subsystem. Use the **rpl maximum** command to configure the maximum number of lines of configuration and number of policies.

The number of lines of configuration includes the beginning and ending statements **,** for example, **route-policy** and **end-policy** . Each line of configuration for sets is also counted.

A line of configuration is counted only once; it is not counted each time it is used. Similarly, any multiple use of policy in an apply statement counts only as one policy.

A user can change the default values for lines and policies but cannot exceed the maximum value, nor can the value for lines and policies be configured lower than the number of lines or policies that are currently configured.

Task ID

Task ID	Operations
route-policy	read, write

Examples

In the following example, the maximum number of RPL system limits are modified:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# rpl maximum lines 50
RP/0/RSP0/CPU0:router(config)# rpl maximum policies 6
```

Related Commands

Command	Description
show rpl maximum, on page 1839	Displays the maximum limits for lines of configuration and number of policies.

rpl set-exit-as-abort

To change the default exit behavior under RPL configuration mode to abort from the RPL configuration mode without saving the configuration, use the **rpl set-exit-as-abort** command in global configuration mode.

rpl set-exit-as-abort

Syntax Description

This command has no keywords or arguments.

Command Modes

Global configuration

Command History

Release	Modification
Release 5.2.0	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.

The default **exit** command acts as end-policy, end-set, or end-if. If the **exit** command is executed under route policy configuration mode, the changes are applied and configuration is updated. This destructs the existing policy. The **rpl set-exit-as-abort** command allows to overwrite the default behavior of the **exit** command under the route policy configuration mode.

Task ID

Task ID	Operations
route-policy	read, write

Examples

This example shows how change the default exit behavior:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# rpl set-exit-as-abort
```

set administrative-distance

To set a route with lower administrative distance such that it is preferred to a route with higher administrative distance, use the **set administrative-distance** command in route policy configuration mode.

set administrative-distance [**number** | **parameter**]

Syntax Description

number	Value assigned to a 8-bit unsigned integer. Range is from 1 to 255.
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 5.2.0	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
route-policy	read, write

Examples

This example shows how to set a route with an administrative value such that it is preferred to a route with higher administrative distance.

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# route-policy sample
RP/0/RSP0/CPU0:router(config-rpl)# set administrative-distance 34
RP/0/RSP0/CPU0:router(config-rpl)# end-policy
RP/0/RSP0/CPU0:router(config-rpl)# exit
RP/0/RSP0/CPU0:router(config)# route bgp 100
RP/0/RSP0/CPU0:router(config-bgp)# address family ipv4 unicast
RP/0/RSP0/CPU0:router(config-bgp-af)# table-policy sample
RP/0/RSP0/CPU0:router(config-bgp-af)# exit
RP/0/RSP0/CPU0:router(config-bgp)# exit
RP/0/RSP0/CPU0:router(config)# end
```

set aigp-metric

To set originating prefixes with accumulated interior gateway protocol (AiGP) attribute information, use the**set aigp-metric** command in route-policy configuration mode.

set aig-metric {**igp-cost***value*}

Syntax Description

igp-cost	Specifies the internal routing protocol cost.
value	Specifies the aigp-metric value. 32- bit decimal number. Range is 0-4294967295.

Command Default

No default behavior or values

Command Modes

Route-policy configuration

Command History

Release	Modification
Release 4.0.0	This command was introduced.
Release 5.0.0	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
route-polcy	read, write

Examples

The following example shows how to set the aigp metric as the igp cost for route-policy aigp_policy:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# route-policy aigp_policy
RP/0/RSP0/CPU0:router(config-rpl)# set aigp-metric igp-cost
```

Related Commands

Commad	Description
aigp	Enables sending and receiving of accumulated interior gateway protocol (AiGP) attribute per neighbor.
aigp send-cost-community	Sends accumulated interior gateway protocol (AiGP) value in cost community.

set community

To set the Border Gateway Protocol (BGP) community attributes in a route, use the **set community** command in route-policy configuration mode.

set community {*community-set-nameinline-community-setparameter*} [**additive**]

Syntax Description

community-set-name	Community set name.
inline-community-set	Inline community set. The inline community set must be enclosed in parentheses.
parameter	Parameter name. The parameter name must be preceded with a "$."
additive	(Optional) Adds communities to communities in the route.

Command Default

No default behavior or values

Command Modes

Route-policy 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **set community** command to set the BGP community attribute.

Note The **set community** command can be used as an action statement within an **if** statement. For a list of all action statements available within an **if** statement, see the **if** command.

Communities are 32-bit values carried in BGP routes. Each route may have zero or more communities in an unordered list.

Use this command to replace the communities in a route or add to them using the optional **additive** keyword.

As with the other community forms that support inline sets, either or both 16-bit portions of the community can be parameterized. Likewise, the names of the well-known communities **internet** (0:0), **no-advertise** (65535:65281), **no-export** (65535:65282), and **local-AS** (65535:65283) can also be used. In an inline community set, each 16-bit portion can also be specified as the **peeras** to express the AS number of the neighbor from which the route was received. If the neighbor AS employs a 4-byte ASN, the IANA-assigned 16-bit value 23456 (AS_TRANS) is used as **peeras** instead.

Without the **additive** keyword, any existing communities (other than the well-known communities) are removed and replaced with the given communities. The **additive** keyword specifies that all communities already present in the route be maintained and the list of communities be added to them.

Task ID

Task ID	Operations
route-policy	read, write

Examples

The following are incomplete configuration examples using the **set community** command:

```
RP/0/RSP0/CPU0:router(config-rpl)# set community (10:24)
RP/0/RSP0/CPU0:router(config-rpl)# set community (10:24, $as:24, $as:$tag)
RP/0/RSP0/CPU0:router(config-rpl)# set community (10:24, internet) additive
RP/0/RSP0/CPU0:router(config-rpl)# set community (10:24, $as:24) additive
RP/0/RSP0/CPU0:router(config-rpl)# set community (10:24, peeras:24) additive
```

set core-tree

To set a Multicast Distribution Tree (MDT) type, use the **set core-tree** command in route-policy configuration mode.

set core-tree {**gre-rosen** | **mldp-inband** | **mldp-partitioned-mp2mp** | **mldp-partitioned-p2mp** | **mldp-rosen** | **rsvp-te-partitioned-p2mp***parameter*}

Syntax Description

gre-rosen	Specifies the IP GRE Rosen core MDT type
mldp-inband	Specifies the MLDP InBand core MDT type
mldp-partitioned-mp2mp	Specifies theMLDP Partitioned MP2MP core MDT type
mldp-partitioned-p2mp	Specifies the MLDP Partitioned P2MP core MDT type
mldp-rosen	Specifies the MLDP Rosen core MDT type
rsvp-te-partitioned-p2mp	Specifies the RSVP TE core core MDT type
parameter	Parameter name. The parameter name must be preceded with a "$."

Command Default

None

Command Modes

Route-policy configuration

Command History

Release	Modification
Release 4.1.0	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
route-policy	read, write

In this example, the Multicast Distribution Tree type is set to IP GRE Rosen core:

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#route-policy policy_mdt_type
RP/0/RSP0/CPU0:router(config-rpl)#set core-tree gre-rosen
```

set dampening

To configure Border Gateway Protocol (BGP) route dampening, use the **set dampening** command in route-policy configuration mode.

set dampening {**halflife** {*minutesparameter*} | **max-suppress** {*minutesparameter*} | **reuse** {*secondsparameter*} | **suppress** {*penalty-unitsparameter*} | **others default**}

Syntax Description

halflife *minutes*	Specifies the time (in minutes) after which a penalty is decreased. After the route has been assigned a penalty, the penalty is decreased by half after the half-life period. The process of reducing the penalty happens every 5 seconds. Range is 1 to 45 minutes.
parameter	Parameter name. The parameter name must be preceded with a "$."
max-suppress *minutes*	Specifies the maximum time (in minutes) a route can be suppressed. Range is 1 to 20000. If the half-life value is allowed to default, the maximum suppress time defaults to 60 minutes.
reuse *seconds*	Unsuppresses a route if the penalty for flapping the route decreases enough to fall below the configured value (in seconds). The process of unsuppressing routes occurs at 10-second increments. Range is 1 to 20000.
suppress *penalty-units*	Specifies a penalty of 1000 each time a route flaps. When a route penalty exceeds the configured limit, it is suppressed. Range is 1 to 20000.
others default	If all four keyword values are not specified in the command, then the command *must* end with **others default** . This designation indicates that any keyword not defined is set to its default.

Command Default

half-life : 15 minutes

max-suppress : 60 minutes (four times the half-life)

reuse : 750 seconds

suppress : 2000 penalty units

Command Modes

Route-policy 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

The BGP protocol supports route dampening using an exponential backoff algorithm. The algorithm is controlled by setting the four supported BGP values: half-life, max-suppress, reuse, and suppress. Use the **set dampening** command to configure BGP route dampening.

Note The **set dampening** command can be used as an action statement within an **if** statement. For a list of all action statements available within an **if** statement, see the **if** command.

A value for at least one of the four keywords must be set. If the **set dampening** command defines values for three or fewer of the supported keywords, then the configuration must end with the **others default** , which indicates that any keyword value not defined in the command is set to its default value.

The keywords may appear in the command in any order.

Task ID

Task ID	Operations
route-policy	read, write

Examples

In the following examples, the half-life is set to 20 minutes and the maximum suppress time is set to

90 minutes. Each command must end with **others default** because three or fewer keywords are defined.

```
RP/0/RSP0/CPU0:router(config-rpl)# set dampening halflife 20 others default
RP/0/RSP0/CPU0:router(config-rpl)# set dampening max-suppress 90 others default
```

In this example, all four keywords are defined, which means the command does not use **others default** .

```
RP/0/RSP0/CPU0:router(config-rpl)# set dampening halflife 15 max-suppress 60 reuse 750
suppress 2000
```

The following command is invalid because it is missing **others default.**

```
RP/0/RSP0/CPU0:router(config-rpl)# set dampening reuse 700
```

In the following example, the parameters are used.

```
RP/0/RSP0/CPU0:router(config-rpl)# set dampening halflife $p1 suppress $p4 reuse $p3
max-suppress $p2
```

set fallback-vrf-lookup

To force forwarding lookup in source VRF for a leaked route, use the **set fallback-vrf-lookup** command in route-policy configuration mode.

set fallback-vrf-lookup

Syntax Description

This command has no keywords or arguments.

Command Default

None

Command Modes

Route-policy configuration

Command History

Release	Modification
Release 6.6.2	This command was introduced.

Usage Guidelines

When L3VPN and EVPN prefixes are imported to a VRF (source VRF) and leaked to an another VRF (destination VRF) with **set fallback-vrf-lookup** command configured, the prefixes are advertised to L3VPN and EVPN peers. When traffic is received from the opposite direction to the destination VRF, first a lookup is performed in the destination VRF, which forces a second lookup in the source VRF. To force this second lookup in the source VRF, use the **set fallback-vrf-lookup** command in route-policy configuration mode.

The **set fallback-vrf-lookup** command has the following route-policy attach points:

- Import from non-default VRF to non-default VRF (use in importing VRF)
- Import from default-VRF to non-default VRF (use in non-default VRF)
- Export to default-VRF from non-default VRF (use in non-default VRF)

Task ID

Task ID	Operation
route-policy	read, write

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# route-policy rpl
RP/0/RSP0/CPU0:router(config-rpl)# set fallback-vrf-lookup
```

set eigrp-metric

To set Enhanced Interior Gateway Routing Protocol (EIGRP) route metrics, use the **set eigrp-metric** command in route-policy configuration mode.

set eigrp-metric *bandwidth delay reliability loading mtu*

Syntax Description

bandwidth	Minimum bandwidth of the route in kilobits per second. Range is 1 to 4294967295.
delay	Route delay in tens of microseconds. Delay is 1 or any positive number that is a multiple of 39.1 nanoseconds. Range is 1 to 4294967295.
reliability	Likelihood of successful packet transmission expressed as a number between 0 and 255. The value 255 means 100 percent reliability; 0 means no reliability.
loading	Effective bandwidth of the route expressed as a number from 1 to 255 (255 is 100 percent loading).
mtu	Minimum maximum transmission unit (MTU) size of the route in bytes. Range is from 1 to 65535.

Command Default

No default behavior or values

Command Modes

Route-policy 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

You can use the **add** command to further offset an existing EIGRP metric value.

Task ID

Task ID	Operations
route-policy	read, write

Examples

In the following example, the EIGRP metrics are adjusted for route policy policy_1:

```
RP/0/RSP0/CPU0:router(config)# route-policy policy_1
RP/0/RSP0/CPU0:router(config-rpl)# set eigrp-metric 1400 120 250 100 1500
RP/0/RSP0/CPU0:router(config-rpl)# end-policy
```

Related Commands

Command	Description
add, on page 1601	Adds an offset to an existing value.

set extcommunity color

To set the Border Gateway Protocol (BGP) color extended community attributes, use the **set extcommunity color** command in route-policy configuration mode.

set extcommunity color *color-value*

Syntax Description

color	Color Extended Community.
color-value	Color Community value.

Command Default

No default behavior or values

Command Modes

Route-policy configuration

Command History

Release	Modification
Release 6.2.1	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.

Use the **set extcommunity color** command to set the color extended community in the overlay prefixes, such as IPv4-unicast SAFI updates between the BGP routers. This command can be used to mark the prefixes with the intended color value.

Note The **set extcommunity color** command can be used as an action statement within an **if** statement. For a list of all action statements available within an **if** statement, see the **if** command.

Task ID

Task ID	Operations
route-policy	read, write

Examples

The following is an incomplete configuration example using the **set extcommunity color** command:

```
RP/0/RSP0/CPU0:router(config-rpl)# set extcommunity color color1
```

set extcommunity cost

To set the Border Gateway Protocol (BGP) cost extended community attributes, use the **set extcommunity cost** command in route-policy configuration mode.

set extcommunity cost {*cost-extcommunity-set-namecost-inline-extcommunity-setparameter*} [**additive**]

Syntax Description

cost-extcommunity-set-name	Cost extended community set name.
cost-inline-extcommunity-set	Inline cost extended community set. The inline cost extended community set must be enclosed in parentheses.
parameter	Parameter name. The parameter name must be preceded with a "$."
additive	(Optional) Adds extended communities for cost to extended communities in the route.

Command Default

No default behavior or values

Command Modes

Route-policy 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **set extcommunity cost** command to either replace the extended communities on the route or add to them using the optional **additive** keyword. Cost community is an extended community used to tie break the best path selection process in BGP so as to have a localized custom decision for packet forwarding. The extended community format defines generic points of insertion (POI) that influence the decision at different points of the bestpath algorithm.



Note The **set extcommunity cost** command can be used as an action statement within an **if** statement. For a list of all action statements available within an **if** statement, see the **if** command.

As with the other extended community forms that support inline sets, either or both portions of the community can be parameterized. Similarly to regular communities, the **additive** keyword can be used to signify adding these extended communities to those that are already present, as opposed to replacing them. Without the **additive** keyword, any existing extended communities for cost (other than the well-known communities) are removed and replaced with the given communities. The **additive** keyword specifies that all extended communities for cost already present in the route be maintained and the set of extended communities be added to them. Well-known communities include internet, local-AS, no-advertise, and no-export.

Task ID

Task ID	Operations
route-policy	read, write

Examples

The following are incomplete configuration examples using the **set extcommunity cost** command:

```
RP/0/RSP0/CPU0:router(config-rpl)# set extcommunity cost (IGP:10:20)
RP/0/RSP0/CPU0:router(config-rpl)# set extcommunity cost (Pre-Bestpath:33:44)
RP/0/RSP0/CPU0:router(config-rpl)# set extcommunity cost (IGP:11:21)
```

set extcommunity rt

To set the Border Gateway Protocol (BGP) route target (RT) extended community attributes, use the **set extcommunity rt** command in route-policy configuration mode.

set extcommunity rt {*rt-extcommunity-set-namert-inline-extcommunity-setparameter*} **additive**

Syntax Description

rt-extcommunity-set-name	Route target extended community set name.
rt-inline-extcommunity-set	Inline route target extended community set. The inline route target extended community set must be enclosed in parentheses.
parameter	Parameter name. The parameter name must be preceded with a "$."
additive	(Optional) Adds extended communities for an RT to extended communities in the route.

Command Default

No default behavior or values

Command Modes

Route-policy 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **set extcommunity rt** command to either replace the extended communities on the route or add to them using the optional **additive** keyword.

Note The **set extcommunity rt** command can be used as an action statement within an **if** statement. For a list of all action statements available within an **if** statement, see the **if** command.

As with the other extended community forms that support inline sets, either or both portions of the community can be parameterized. Similarly to regular communities, the **additive** keyword can be used to signify adding these extended communities to those that are already present, as opposed to replacing them.

Task ID

Task ID	Operations
route-policy	read, write

Examples

The following are incomplete configuration examples using the **set extcommunity rt** command:

```
RP/0/RSP0/CPU0:router(config-rpl)# set extcommunity rt (10:24)
RP/0/RSP0/CPU0:router(config-rpl)# set extcommunity rt (10:24, $as:24, $as:$tag)
RP/0/RSP0/CPU0:router(config-rpl)# set extcommunity rt (10:24, internet) additive
RP/0/RSP0/CPU0:router(config-rpl)# set extcommunity rt (10:24, $as:24) additive
```

Without the **additive** keyword, any existing extended communities for cost (other than the well-known communities) are removed and replaced with the given communities. The **additive** keyword specifies that all extended communities for cost already present in the route be maintained and the list of extended communities be added to them.

set ip-precedence

To set the IP precedence, use the **set ip-precedence** command in route-policy configuration mode.

set ip-precedence {*numberparameter*}

Syntax Description

number	Value of the precedence. The precedence value can be a number from 0 to 7: • **7** —network (set packets with network control precedence) • **6** —internet (set packets with internetwork control precedence) • **5** —critical (set packets with critical precedence) • **4** —flash-override (set packets with flash override precedence) • **3** —flash (set packets with flash precedence) • **2** —immediate (set packets with immediate precedence) • **1** —priority (set packets with priority precedence) • **0** —routine (set packets with routine precedence)
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 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **set ip-precedence** command to set the IP precedence to classify packets. This command is supported at the BGP table-policy attachpoint. Prefixes are marked for subsequent processing in the forwarding plane. After QoS Policy Propagation through Border Gateway Protocol (BGP) (QPPB) is enabled on an interface, corresponding traffic shaping and policing is completed using packet classification based on the IP precedence or QoS group ID. See *Modular QoS Configuration Guide for Cisco ASR 9000 Series Routers* for information on QPPB.

Task ID

Task ID	Operations
route-policy	read, write

Examples

This example shows how use **set ip-precedence** command:

```
RP/0/RSP0/CPU0:router(config)# route-policy policy_1
RP/0/RSP0/CPU0:router(config-rpl)# set ip-precedence 3
```

```
RP/0/RSP0/CPU0:router(config-rpl)# end-policy
```

set isis-metric

To set the Intermediate System-to-Intermediate System (IS-IS) metric attribute value, use the **set is-is metric** command in route-policy configuration mode.

set isis-metric {*numberparameter*}

Syntax Description

number	24-bit integer number. Range is from 0 to 16777215.
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 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **set isis-metric** command to set the IS-IS metric attribute value for routes that are redistributed into IS-IS.

Task ID

Task ID	Operations
route-policy	read, write

Examples

In the following example, the IS-IS metric attribute value is set to 1000:

```
RP/0/RSP0/CPU0:router(config)# route-policy policy_1
RP/0/RSP0/CPU0:router(config-rpl)# set isis-metric 1000
RP/0/RSP0/CPU0:router(config-rpl)# end-policy
```

set label

To set the Border Gateway Protocol (BGP) label attribute value, use the **set label** command in route-policy configuration mode.

set label {**explicit-null** | **implicit-null***parameter*}

Syntax Description

explicit-null	Sets the label to the well-known explicit value of 0.
implicit-null	Sets the label to the well-known implicit value of 3.
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 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **set label** command in a route policy at the allocate label attachpoint to set the label to explicit-null or implicit-null based on deployment preference. During inter-AS operation, the ASBR sends some of its own loopbacks to other its peers and labels them either implicit null or explicit null.

Examples

The following example shows how to set the labels:

```
RP/0/RSP0/CPU0:router(config)# route-policy labelpolicy
RP/0/RSP0/CPU0:router(config-rpl)# if destination in (206.141.1.0/24) then
RP/0/RSP0/CPU0:router(config-rpl)# set label explicit-null
RP/0/RSP0/CPU0:router(config-rpl)# elseif destination in (206.141.3.0/24) then
RP/0/RSP0/CPU0:router(config-rpl)# drop
RP/0/RSP0/CPU0:router(config-rpl)# elseif destination in (206.141.4.0/24) then
RP/0/RSP0/CPU0:router(config-rpl)# set label explicit-null
RP/0/RSP0/CPU0:router(config-rpl)# endif
RP/0/RSP0/CPU0:router(config-rpl)# end-policy
```

set label-mode

To set the type of Border Gateway Protocol (BGP) label mode, use the **set label-mode** command in route-policy configuration mode. This command does not have a **no** form.

set label-mode { **per-ce** | **per-vrf** | **per-prefix** }

Syntax Description

per-ce	Specifies that the same label is used for all routes advertised from a unique customer edge (CE) peer or router.
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.

Command Default

Per-prefix label mode.

If a policy attached at label-mode attachpoint evaluates to pass and a **label mode** is not explicitly set, **per-prefix** is used as a default label mode.

Command Modes

Route-policy configuration

Command History

Release	Modification
Release 4.3.1	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.

Use the **set label-mode** command in a route policy at the label-mode attachpoint to set the type of label mode to per-ce or per-vrf or per-prefix, based on deployment preference.

Per-vrf label mode is not supported for Carrier Supporting Carrier (CSC) network with internal and external BGP multipath setup.

Task ID

Task ID	Operation
route-policy	read, write

This example shows how to set the type of label-mode to per-ce:

```
Router(config)# route-policy set_label_mode
Router(config-rpl)# set label-mode per-ce
Router(config-rpl)# end-policy
```

This example shows how to set the type of label-mode to per-vrf:

```
Router(config)# route-policy set_label_mode
```

```
Router(config-rpl)# set label-mode per-vrf
Router(config-rpl)# end-policy
```

This example shows how to set the type of label-mode to per-prefix:

```
Router(config)# route-policy set_label_mode
Router(config-rpl)# set label-mode per-prefix
Router(config-rpl)# end-policy
```

Related Commands

Command	Description
route-policy (RPL)	Defines a route policy and enters route-policy configuration mode.

set large-community

To set the Border Gateway Protocol (BGP) large-community attributes in a route, use the **set large-community** command in route-policy configuration mode.

set large-community { *large-community-set-name inline-large-community-set parameter* } [**additive**]

Syntax Description

large-community-set-name	Large-community set name.
inline-large-community-set	Inline large-community set. The inline large-community set must be enclosed in parentheses.
parameter	Parameter name. The parameter name must be preceded with a "$."
additive	(Optional) Adds large-communities to large-communities in the route.

Command Default

No default behavior or values

Command Modes

Route-policy configuration

Command History

Release	Modification
Release 6.3.1	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.

The large communities 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.

In route-policy statements, each integer in the BGP large community can be replaced by the following expression:

- peeras — This expression is replaced by the AS number of the neigbhor from which the community is received or to which the community is sent, as appropriate.

Note The **set large-community** command can be used as an action statement within an **if** statement. For a list of all action statements available within an **if** statement, see the **if** command.

Without the **additive** keyword, any existing large communities are removed and replaced with the given large communities. The **additive** keyword specifies that all communities already present in the route be maintained and the list of communities be added to them. However the **additive** keyword does not result in duplicate entries. If a particular large community is attached to a route and you specify the same large community again with the **additive** keyword in the set statement, then the specified large community is not added again. The merging operation removes duplicate entries. This also applies to the **peeras** keyword.

Task ID

Task ID	Operations
route-policy	read, write

The peeras expression in this example is replaced by the AS number of the neighbor from which the BGP large community is received or to which the community is sent, as appropriate.

In this example, if the route-policy mordac is applied to a neighbor, the ASN of which is 1, then the large community (1:2:3) is set only once.

```
RP/0/RP/0/RSP0/CPU0:router#config
RP/0/RSP0/CPU0:router(config)#route-policy mordac
RP/0/RSP0/CPU0:router(config-rpl)#set large-community (1:2:3, peeras:2:3)
RP/0/RSP0/CPU0:router(config-rpl)#end-set
RP/0/RSP0/CPU0:router(config)#large-community-set catbert
RP/0/RSP0/CPU0:router(config-largecomm)#1:2:3,
RP/0/RSP0/CPU0:router(config-largecomm)#5:2:3
RP/0/RSP0/CPU0:router(config-largecomm)#end-set
RP/0/RSP0/CPU0:router(config)#route-policy wally
RP/0/RSP0/CPU0:router(config-rpl)#set large-community catbert additive
RP/0/RSP0/CPU0:router(config-rpl)#end-set
```



Note You should configure the **send-community-ebgp** command to send large communities to ebgp neighbors.

set level

To configure the Intermediate System-to-Intermediate System (IS-IS) link-state packet (LSP) level advertised to redistributed routes, use the **set level** command in route-policy configuration mode.

set level {**level-1** | **level-2** | **level-1-2***parameter*}

Syntax Description

level-1	Specifies that redistributed routes are advertised in the Level 1 LSP of the router.
level-2	Specifies that redistributed routes are advertised in the Level 2 LSP of the router.
level-1-2	Specifies that redistributed routes are advertised in Level 1 and Level 2 LSPs of the router.
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 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the IS-IS **set level** command to configure the LSP level advertised to redistributed routes.

Note The **set level** command can be used as an action statement within an **if** statement. For a list of all action statements available within an **if** statement, see the **if** command.

This command supports parameterization of the **level** keyword.

Task ID

Task ID	Operations
route-policy	read, write

Examples

In the following example, the level is set to Level 2:

```
RP/0/RSP0/CPU0:router(config)# route-policy bgp_isis_redist
RP/0/RSP0/CPU0:router(config-rpl)# if destination in (172.2.0.0/16 ge 16) then
RP/0/RSP0/CPU0:router(config-rpl)# set level level-2
RP/0/RSP0/CPU0:router(config-rpl)# endif
RP/0/RSP0/CPU0:router(config-rpl)# end-policy
```

set local-preference

To set the Border Gateway Protocol (BGP) local preference attribute in a route, use the **set local-preference** command in route-policy configuration mode.

set local-preference {*numberparameter*}

Syntax Description

number	Value assigned to a 32-bit unsigned integer. Range is 0 to 4294967295.
parameter	Parameter name. The parameter name must be preceded with a "$."

Command Default

Default value is 100.

Command Modes

Route-policy 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **set local-preference** command to specify a preference value for the autonomous system path. Local preference is a nontransitive (does not cross autonomous system boundaries) attribute and is the second metric considered in the BGP best path calculation (the highest local preference is chosen). Weight is the first metric evaluated for best path, but it is local to the router and propagates only to iBGP peers. See the *Implementing BGP* module of the *Routing Configuration Guide for Cisco ASR 9000 Series Routers* for information on the BGP best path calculation.

Note The **set local-preference** command can be used as an action statement within an **if** statement. For a list of all action statements available within an **if** statement, see the **if** command.

The local preference is a 32-bit unsigned integer.

Task ID

Task ID	Operations
route-policy	read, write

Examples

In the following example, the local preference value is set to 10:

```
RP/0/RSP0/CPU0:router(config-rpl)# set local-preference 10
```

set med

To set the Border Gateway Protocol (BGP) Multi Exit Discriminator (MED) attribute, use the **set med** command in route-policy configuration mode.

set med {*numberparameter* | **igp-cost** | {**+** | {*numberparameter*} | **-** | {*numberparameter*}} | **max-reachable**}

Syntax Description

number	Value assigned to a 32-bit unsigned integer. Range is 0 to 4294967295.
parameter	Parameter name. The parameter name must be preceded with a "$."
igp-cost	Sets the MED value to the cost for the Interior Gateway Protocol (IGP) route to resolve the next-hop of the BGP route.
+ | **-**	Sets the MED to the MED plus or minus a static offset. An integer or parameter must follow the plus or minus.
max-reachable	Sets the MED value to the maximum possible value of 4294967295.

Command Default

No default behavior or values

Command Modes

Route-policy 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **set med** command to set the MED value, which is a 32-bit unsigned integer.

Note The **set med** command can be used as an action statement within an **if** statement. For a list of all action statements available within an **if** statement, see the **if** command.

This command can take the following as argument values: an integer, a parameter, the **igp-cost** keyword, or a mathematical operator (either plus or minus) followed by an integer or a parameter. Setting the MED to the IGP cost is supported on outbound BGP policies only. The MED cannot be set to the IGP cost in policies applied to other BGP attach points.

The **max-reachable** keyword sets the MED to the maximum value while leaving the route reachable.

The plus or minus variants allow the user to set the MED to the MED plus or minus a static offset. The variants that allow a user to add or subtract offsets to the MED value are also range checked for underflow or overflow. If the value underflows as a result of subtraction, then the MED value is set to zero. If the value overflows, the value is set to 4294967295, which is the maximum value for MED. when MED is set to 4294967295, the route is unreachable.

Task ID

Task ID	Operations
route-policy	read, write

Examples

The following two examples show how to set the MED to a value that is either specified directly (using the integer 156) or passed to the policy as a parameter:

```
RP/0/RSP0/CPU0:router(config-rpl)# set med 156
RP/0/RSP0/CPU0:router(config-rpl)# set med $med_param
```

The following example shows how to instruct BGP to automatically set the MED value to the cost of the IGP route that resolves the next-hop of the BGP route:

```
RP/0/RSP0/CPU0:router(config-rpl)# set med igp-cost
```

set metric-type (IS-IS)

To configure the integrated Intermediate System-to-Intermediate System (IS-IS) metric type, use the **set metric-type** command in route-policy configuration mode.

set metric-type {**internal** | **external***parameter*}

Syntax Description

internal	Sets metric type to internal.
external	Sets the metric type to external.
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 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the IS-IS **set metric-type** command to control whether IS-IS treats the metric as an internal or external metric.



Note The **set metric-type** command can be used as an action statement within an **if** statement. For a list of all action statements available within an **if** statement, see the **if** command.

This command does not support parameterization.

Task ID

Task ID	Operations
route-policy	read, write

Examples

In the following example, the IS-IS metric type is set to internal:

```
RP/0/RSP0/CPU0:router(config-rpl)# set metric-type internal
```

set metric-type (OSPF)

To control how Open Shortest Path First (OSPF) computes the cost for a route, use the **set metric-type** command in route-policy configuration mode.

set metric-type {**type-1** | **type-2***parameter*}

Syntax Description

type-1	Uses the cost set on the route plus the topology-related costs in the calculation for Type 1 metrics.
type-2	Uses only the cost set on the route in the calculation for Type 2 metrics.
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 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the OSPF **set metric-type** command to control whether OSPF treats the cost as a Type 1 or Type 2 metric.

Note The **set metric-type** command can be used as an action statement within an **if** statement. For a list of all action statements available within an **if** statement, see the **if** command.

The value of Type 1 or Type 2 controls how OSPF computes the cost for this route. For Type 2 metrics, only the cost set on the route is used. For Type 1 metrics, the cost set on the route plus the topology- related costs are used in the calculation.

This command does not support parameterization.

Task ID

Task ID	Operations
route-policy	read, write

Examples

In the following example, the OSPF metric type is set to Type 1:

```
RP/0/RSP0/CPU0:router(config-rpl)# set metric-type type-1
```

set next-hop

To replace the next-hop associated with a given route, use the **set next-hop** command in route-policy configuration mode.

set next-hop {*ipv4-address [destination-vrf]*| *ipv6-address [destination-vrf]* | **discard** | **ipv6-global** *parameter* | **peer-address** | **self**}

Syntax Description

ipv4-address	Valid IPv4 address.
ipv6-address	Valid IPv6 address.
destination-vrf	(Optional) Specifies that the next-hop of the route should be resolved in destination VRF context. This keyword is available when an IPv4 or IPv6 address or parameter is used.
discard	Sets next-hop as Null0 interface.
parameter	Parameter name. The parameter name must be preceded with a "$."
peer-address	Sets the next-hop to the IP address of the remote Border Gateway Protocol (BGP) peer.
self	Sets itself as the next-hop.
unchanged	Sets next-hop unchanged

Command Default

No default behavior or values

Command Modes

Route-policy configuration

Command History

Release	Modification
Release 3.7.2	This command was introduced.
Release 4.3.0	The **discard** keyword was added.

Usage Guidelines

Use the **set next-hop** command to replace the next-hop associated with a specific address.

The next hop destination is selected according to the address family. Example: for ipv4 address-family, the IPv4 address is used and for ipv6 address-family, the IPv6 address is used.

Note The **set next-hop** command can be used as an action statement within an **if** statement. For a list of all action statements available within an **if** statement, see the **if** command.

Use the **set next-hop peer-address** command to set the next-hop to the address of the BGP neighbor, where this policy is attached.

The next-hop is a valid IPv4 address entered as a dotted decimal or an IPv6 address entered as a colon-separated hexadecimal.

It is not possible to use this command to set the BGP IPv6 link-local next-hop.

The **destination-vrf** keyword is used mainly in Layer 3 VPN networks when importing routes.

The **set next-hop discard** configuration is used in the neighbor inbound policy. When this config is applied to a path, the primary next-hop is still be associated with the actual path but the RIB is updated with next-hop set to Null0. Even if the primary received nexthop is unreachable, the Remotely Triggered Blackhole (RTBH) path will be considered reachable and will be a candidate in the bestpath selection process. The RTBH path is readvertised to other peers with either the received next-hop or nexthop-self based on normal BGP advertisement rules.

Task ID

Task ID	Operations
route-policy	read, write

Examples

In the following example, the next-hop is set to a valid IPv4 address:

```
RP/0/RSP0/CPU0:router(config-rpl)# set next-hop 10.0.0.5
```

In this example, the next-hop is set to a parameter value $nexthop:

```
RP/0/RSP0/CPU0:router(config-rpl)# set next-hop $nexthop
```

In this example, the next-hop is set to a valid IPv4 address with a destination VRF context:

```
RP/0/RSP0/CPU0:router(config-rpl)# set next-hop 10.0.0.5 destination-vrf
```

set origin

To change the Border Gateway Protocol (BGP) origin attribute, use the **set origin** command in route-policy configuration mode.

set origin {**igp** | **incomplete** | **egp***parameter*}

Syntax Description

igp	Sets the origin type to Interior Gateway Protocol (IGP).
incomplete	Sets the origin type to incomplete.
egp	Sets the origin type to Exterior Gateway Protocol (EGP).
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 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **set origin** command to change the origin attribute.

Note The **set origin** command can be used as an action statement within an **if** statement. For a list of all action statements available within an **if** statement, see the **if** command.

The origin of a Border Gateway Protocol (BGP) route is **igp** , **egp** , or **incomplete** .

Task ID

Task ID	Operations
route-policy	read, write

Examples

In the following example, the origin attribute is set to EGP:

```
RP/0/RSP0/CPU0:router(config-rpl)# set origin egp
```

set ospf-metric

To set an Open Shortest Path First (OSPF) protocol metric attribute value, use the **set ospf-metric** command in route-policy configuration mode.

set ospf-metric {*numberparameter*}

Syntax Description

number	Value assigned to a 24-bit unsigned integer. Range is 0 to 4294967295.
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 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **set ospf-metric** command to set the metric for routes that are redistributed into OSPF. The OSPF metric operator accepts either an integer value or a parameter.

Task ID

Task ID	Operations
route-policy	read, write

Examples

In the following example, the OSPF metric attribute value is set to 1000:

```
RP/0/RSP0/CPU0:router(config)# route-policy policy_1
RP/0/RSP0/CPU0:router(config-rpl)# set ospf-metric 1000
RP/0/RSP0/CPU0:router(config-rpl)# end-policy
```

set path-selection

Sets Border Gateway Protocol (BGP) path selection criteria.

To set path selection criteria and install or advertise the path for the Border Gateway Protocol, use the **set path-selection** command in route-policy configuration mode.

set path-selection { **backup** *number* | **group-best** | **all** | **best-path** } **[install [multipath-protect]]** [**advertise**] **[multipath-protect-advertise]**

Syntax Description

backup	Specifies the BGP backup path.
number	Specifies the BGP backup path number, which must be a value of 1.
group-best	Specifies all the BGP group best paths. A group best path is the best path received from an AS. For example, if paths are received from 3 ASes, then there are 3 group best paths.
all	Specifies all BGP paths.
best-path	Specifies the BGP best path.
install	Installs the path.
multipath-protect	Selects a backup path that is not also a multipath. All selected paths should have unique next hops. For example, if two paths have the same next hop, then one of them is not selected. The backup path is selected such that it does not have the same next hop as any other path that is a multipath. **Note** Multipaths are configured with the **maximum-paths** command.
advertise	Advertises the path.
multipath-protect-advertise	Installs and advertises the multipath protection along the multipath.
multipath-protect-advertise	**Note** Use this keyword when you do not need to configure the backup path selection.

Command Default

None

Command Modes

Route-policy configuration

Command History

Release	Modification
Release 4.0.0	This command was introduced.

Release	Modification
Release 4.0.1	The **multipath-protect** keyword was added.
Release 7.3.1	The **multipath-protect-advertise** keyword was added

Usage Guidelines

This command is used with the **additional-path selection** command within an appropriate route-policy to calculate backup paths and to enable Prefix Independent Convergence (PIC) functionality. See the *BGP Prefix Independent Convergence Unipath Primary/Backup* section in *Routing Configuration Guide for Cisco ASR 9000 Series Routers* for details on the PIC functionality.

The group-best is the set of paths that are the best paths from the paths of the same autonomous system (AS). All the paths that are selected as the group-best set should be advertised to peers, however, if there are multiple paths that are selected as group-best because they come from different ASs but having the same next-hop the XR router advertises only one path from these paths to the peer device.

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.

Examples

The following example shows how to set the backup path for a route-policy:

```
Router# configure
Router(config)# route-policy path_selection_policy
Router(config-rpl)# set path-selection backup 1 advertise
```

The following example shows how to set the backup path for a route-policy and how to installs and advertise the multipath protection along the multipath:

```
Router# configure
Router(config)# route-policy path_selection_policy
Router(config-rpl)# set path-selection backup 1 multipathprotect-advertise
```

Related Commands

Command	Description
additional-paths selection	Configures additional paths selection capability for a prefix
additional-paths send	Configures send capability of multiple paths for a prefix to the capable peers,
additional-paths receive	Configures receive capability of multiple paths for a prefix to the capable peers.
advertise best-external	Advertises the best–external path to the iBGP and route-reflector peers,

set qos-group (RPL)

To set the quality of service (QoS) group, use the **set qos-group** command in route-policy configuration mode:

set qos-group {*numberparameter*}

Syntax Description

number	QoS group ID. Range is from 0 to 31.
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 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **set qos-group** command to set the QoS group to classify packets.

This command is supported at the BGP table-policy attachpoint. Prefixes are marked for subsequent processing in the forwarding plane. After QoS Policy Propagation through Border Gateway Protocol (BGP) (QPPB) is enabled on an interface, corresponding traffic shaping and policing is completed using packet classification based on the IP precedence or QoS group ID. See the *Modular QoS Configuration Guide for Cisco ASR 9000 Series Routers* for information on QPPB.

Task ID

Task ID	Operations
route-policy	read, write

Examples

This example shows how to use **set qos-group** command:

```
RP/0/RSP0/CPU0:router(config)# route-policy policy_1
RP/0/RSP0/CPU0:router(config-rpl)# set qos-group 12
RP/0/RSP0/CPU0:router(config-rpl)# end-policy
```

set rib-metric

To set the Routing Information Base (RIB) metric attribute value for a table policy, use the **set rib-metric** command in route-policy configuration mode:

set rib-metric {*numberparameter*}

Syntax Description

number	Value assigned to a 32-bit unsigned integer. Range is 0 to 4294967295.
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 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **set rib-metric** command set the RIB metric attribute value for BGP routes.

Every route in the RIB has a metric associated with it, signifying the cost to reach a specific destination based on link characteristics. The **set rib-metric** command modifies the RIB metric while installing BGP routes into RIB, enabling the upgrading or downgrading of the BGP route installed in RIB.

Task ID

Task ID	Operations
route-policy	read, write

Examples

In the following example, the RIB metric attribute is set to 1000:

```
RP/0/RSP0/CPU0:router(config)# route-policy policy_1
RP/0/RSP0/CPU0:router(config-rpl)# set rib-metric 1000
RP/0/RSP0/CPU0:router(config-rpl)# end-policy
```

set rip-metric

To set Routing Information Protocol (RIP) metric attributes, use the **set rip-metric** command in route-policy configuration mode.

set rip-metric {*numberparameter*}

Syntax Description

number	Value assigned to a 4-bit unsigned integer. Range is from 0 to 16.
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 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **set rip-metric** command to set the cost attribute for routes that are redistributed into RIP.

You can use the **add** command to increment the RIP metric value.

Task ID

Task ID	Operations
route-policy	read, write

Examples

In the following example, the RIP metric number is adjusted for route policy policy_1:

```
RP/0/RSP0/CPU0:router(config)# route-policy policy_1
RP/0/RSP0/CPU0:router(config-rpl)# set rip-metric 10
RP/0/RSP0/CPU0:router(config-rpl)# end-policy
```

Related Commands

Command	Description
add, on page 1601	Adds an offset to an existing value.

set rip-tag

To set a route tag attribute for Routing Information Protocol (RIP) routes, use the **set rip-tag** command in route-policy configuration mode.

set rip-tag {*numberparameter*}

Syntax Description

number	Value assigned to a 16-bit unsigned integer. Range is from 0 to 65535.
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 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **set rip-tag** command to set the RIP tag attribute for routes that are redistributed into RIP. The RIP tag operator accepts either an integer value or a parameter.

Task ID

Task ID	Operations
route-policy	read, write

Examples

In the following example, the RIP tag is adjusted for route policy policy_1:

```
RP/0/RSP0/CPU0:router(config)# route-policy policy_1
RP/0/RSP0/CPU0:router(config-rpl)# set rip-tag 1000
RP/0/RSP0/CPU0:router(config-rpl)# end-policy
```

set rpf-topology

To set reverse-path forwarding (RPF) to any default or nondefault tables for particular sources and groups, use the **set rpf-topology** command in routing policy configuration mode.

set rpf-topology [**vrf** *vrf-name*] {**ipv4** | **ipv6**} {**unicast** | **multicast***parameter*} **topology** *table-name*

Syntax Description

ipv4	[Optional] Specifies IPv4 address prefixes.
ipv6	[Optional] Specifies IPv6 address prefixes.
unicast	Specifies unicast address prefixes.
multicast	Specifies multicast address prefixes.
parameter	Parameter name. The parameter name must be preceded with a "$."
topology	Specifies the default or non-default topology table for the source or group.
table-name	Alphanumeric name string.

Command Default

Default or current topology setting.

Command Modes

Routing policy 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 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
multicast	read

Examples

The following example shows how to execute the **set rpf-topology** command:

```
RP/0/RSP0/CPU0:router# config
RP/0/RSP0/CPU0:router(config)# route-policy green
RP/0/RSP0/CPU0:router(config-rpl)# set rpf-topology ipv6 multicast topology t12
```

The following example shows the use of **set rpf-topology** command in the context of creating an RPF for a topology table in multiple topologies:

```
route-policy mt4-p1
  if destination in (225.0.0.1, 225.0.0.11) then
```

```
    set rpf-topology ipv4 multicast topology t201
  elseif destination in (225.0.0.2, 225.0.0.12) then
    set rpf-topology ipv4 multicast topology t202
  elseif destination in (225.0.0.3, 225.0.0.13) then
    pass
  endif
end-policy
!


route-policy mt4-p3
  if destination in (225.0.0.8) then
    set rpf-topology ipv4 multicast topology t208
  elseif destination in (225.0.0.9) then
    set rpf-topology ipv4 multicast topology t209
  elseif destination in (225.0.0.10) then
    set rpf-topology ipv4 multicast topology t210
  else
    drop
  endif
end-policy
!
```

Related Commands

Command	Description
rpf topology	Assigns a route policy in PIM to select a reverse-path forwarding (RPF) topology.

set rtset route-limit

To set limits on paths in the inbound neighbor policy in BGP, particularly when acting as a route-reflector, use the **set rtset route-limit** command in the global configuration mode. If such a path is accepted, BGP adds a flag to the path, BPATH_RTSET_NET_COUNT, to indicate that the path is subjected to the limit.

To remove a limit, use the **no** form of this command.

set rtset route-limit

Syntax Description

limit-value	Displays the 32-unit quantity.

Command Default

No default behavior or values.

Command Modes

Global configuration

Command History

Release	Modification
Release 5.0	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
route-policy	read

Examples

Based on service agreement, if a user VPN1 advertises VPN prefixes with RT-set {1:1, 2:2} and user VPN2 advertises VPN prefixes with RTs {3:3, 4:4}. On the DUT, the user wants to limit VPN1 prefixes to 100, VPN2 prefixes to 50, and for all other users, limit each of their prefixes to 80. Note that the fallback limit is per user, not a cumulative one. To achieve this, the user must configure the following route-policy:

```
route-policy RTSET-LIMIT
  if extcommunity rt matches-every (1:1, 2:2) then
     set rtset-route-limit 100
  elseif extcommunity rt matches-every (3:3, 4:4) then
     set reset-route-limit 50
  else
     set reset-route-limit 80
  endif
end-policy

With this configuration, the RR will keep
   (i) at most 100 prefixes that have a path with RT-set containing
       {1:1, 2:2},
  (ii) at most 50 prefixes that have a path with RT-set containing {3:3, 4:4},
       and
 (iii) at most 80 prefixes that only have paths with RT-set that do not
       contain either {1:1, 2:2} or {3:3, 4:4}
```

set spf-priority

To set OSPF Shortest Path First (SPF) priority, use the set spf-priority command in route-policy configuration mode.

set spf-priority {**critical** | **high** | **medium**}

Syntax Description

critical	Sets critical priority for SPF
high	Sets high priority for SPF
medium	Sets medium priority for SPF

Command Default

None

Command Modes

Route-policy configuration

Command History

Release	Modification
Release 3.7.0	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
route-policy	read, write

This example sets SPF priority as critical:

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#route-policy policy_spf_priority
RP/0/RSP0/CPU0:router(config-rpl)#set spf-priority critical
```

Related Commands

Command	Description
spf prefix-priority (OSPF)	Prioritizes OSPFv2 prefix installation into the global Routing Information Base (RIB) during Shortest Path First (SPF) run.

set tag

To set the tag attribute, use the **set tag** command in route-policy configuration mode.

set tag {*numberparameter*}

Syntax Description

number	Value assigned to a 32-bit unsigned integer. Range is from 0 to 4294967295.
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 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **set tag** command to set the tag attribute.

Note The **set tag** command can be used as an action statement within an **if** statement. For a list of all action statements available within an **if** statement, see the **if** command.

Tags are routing-protocol independent 32-bit integers that can be associated with a given route in the Routing Information Base (RIB).

For the Border Gateway Protocol (BGP), the tag attribute can be set only at the table-policy attach point.

Task ID

Task ID	Operations
route-policy	read, write

Examples

In the following example, the tag attribute is set to 10:

```
RP/0/RSP0/CPU0:router(config-rpl)# set tag 10
```

In this example, the tag attribute is set to a parameter value $tag_param:

```
RP/0/RSP0/CPU0:router(config-rpl)# set tag $tag_param
```

set traffic-index

To set the traffic index attribute, use the **set traffic-index** command in route-policy configuration mode.

set traffic-index {*numberparameter* | **ignore**}

Syntax Description

number	Integer value assigned to the traffic index attribute. Range is 1 to 63.
parameter	Parameter name. The parameter name must be preceded with a "$."
ignore	Specifies that Border Gateway Protocol (BGP) policy accounting is not done.

Command Default

No default behavior or values

Command Modes

Route-policy 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **set traffic-index** command to set the traffic index attribute.

Note The **set traffic-index** command can be used as an action statement within an **if** statement. For a list of all action statements available within an **if** statement, see the **if** command.

Traffic index is a special attribute for BGP. It is used as an index to a set of counters that are maintained by forwarding hardware. It is also used to track packet and byte counters that are forwarded using routes with specific attributes. These counters can be enabled and disabled on an individual interface basis.

The traffic index attribute can be set only at the table-policy attach point, and can take a value from 1 to 63, or a value of **ignore** . If the traffic index is set to **ignore** , then BGP policy accounting is not done. Parameterization of this value is also supported.

Task ID

Task ID	Operations
route-policy	read, write

Examples

In the following example, a policy is created in which the traffic index is set to 10 for all routes that originated in autonomous system 1234:

```
RP/0/RSP0RP0/CPU0:router(config)# route-policy count-as-1234
```

```
RP/0/RSP0RP0/CPU0:router(config-rpl)# if as-path originates-from '1234' then
RP/0/RSP0RP0/CPU0:router(config-rpl-if)# set traffic-index 10
RP/0/RSP0RP0/CPU0:router(config-rpl-if)# else
RP/0/RSP0RP0/CPU0:router(config-rpl-if)# pass
RP/0/RSP0RP0/CPU0:router(config-rpl-if)# endif
RP/0/RSP0RP0/CPU0:router(config-rpl)# end-policy
```

This policy could then be attached using the BGP **table-policy** command. The counters could then be enabled on various interfaces with the appropriate commands.

set vpn-distinguisher

To change the Border Gateway Protocol (BGP) VPN distinguisher attribute, use the **set vpn-distinguisher** command in route-policy configuration mode.

set vpn-distinguisher {*numberparameter*}

Syntax Description

number	Value assigned to a 32-bit unsigned integer. Range is from 1 to 4294967295.
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 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **set vpn-distinguisher** command to change the VPN distinguisher attribute.

Note The **set origin** command can be used as an action statement within an **if** statement. For a list of all action statements available within an **if** statement, see the **if** command.

A VPN distinguisher is used in Layer 3 VPN networks for enhanced individual VPN control and to avoid route target mapping at AS boundaries in inter-AS VPN networks. Route target extended communities are removed at neighbor outbound, and the VPN distinguisher value is applied on the BGP route as an extended community. When the route is received on a neighboring router in another AS, the VPN distinguisher is removed and mapped to a route target extended community.

Task ID

Task ID	Operations
route-policy	read, write

Examples

In the following example, the VPN distinguisher attribute is set to 456:

```
RP/0/RSP0/CPU0:router(config-rpl)# set vpn-distinguisher 456
```

set weight

To set the weight value for Border Gateway Protocol (BGP) routes, use the **set weight** command in route-policy configuration mode.

set weight {*numberparameter*}

Syntax Description

number	Number assigned to the weight value for BGP routes. Weight is 16 bits. Range is 0 to 65535.
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 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **set weight** command to set the weight value for BGP routes.

Note The **set weight** command can be used as an action statement within an **if** statement. For a list of all action statements available within an **if** statement, see the **if** command.

A weight is a value that can be applied to a route to override the BGP local preference. This is not a BGP attribute announced to BGP peer routers. RPL can be used to set the weight value.

Given two BGP routes with the same network layer reachability information (NLRI), a route with a higher weight is selected, no matter what the values of other BGP attributes may be. However, weight only has significance on the local router. It is not sent from one BGP speaker to another, even within the same autonomous system.

On Cisco routers, if a BGP route is sourced by the local router, its weight is automatically set to 32768; if the BGP route is learned from another router, its weight is automatically set to 0. Thus, by default, locally sourced routes are preferred over BGP learned routes.

Task ID

Task ID	Operations
route-policy	read, write

Examples

In the following example, the weight of the route is set to 10 and then to a parameter value $weight_param:

```
RP/0/RSP0/CPU0:router(config-rpl)# set weight 10
RP/0/RSP0/CPU0:router(config-rpl)# set weight $weight_param
```

show rpl

To display system-wide RPL configuration, use the **show rpl** command in EXEC mode.

show [**running-config**] **rpl** [**maximum** {**lines** *configuration-limit* | **policies** *policies-limit*} | **editor** {**emacs** | **nano** | **vim**}]

Syntax Description

running-config	(Optional) Displays configuration-limit argument.
maximum	(Optional) Displays the maximum number of lines of configuration and number of policies.
lines *configuration-limit*	(Optional) Displays the number of lines to which configuration is limited. Range is 1 to 131072. The *configuration-limit* argument is available if the **running-config** keyword is specified.
policies *policies-limit*	(Optional) Displays the limit on the number of policies. Range is 1 to 5000. The *configuration-limit* argument is available if the **running-config** keyword is specified.
editor	(Optional) Specifies the default RPL editor. This keyword is available if the **running-config** keyword is specified.
emacs	(Optional) Displays the default RPL editor to Micro Emacs.
nano	(Optional) Displays the default RPL editor to nano.
vim	(Optional) Displays the default RPL editor to Vim.

Command Default

No default behavior or values

Command Modes

EXEC

Command History

Release	Modification
Release 3.7.2	This command was introduced.
Release 7.11.1	The **Nano** and **Emacs** keyword was deprecated.

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
route-policy	read, write

Examples

The following shows the output of the **show running-config rpl** command:

```
RP/0/RSP0/CPU0:router# show running-config rpl

extcommunity-set rt ext_comm_set_rt_ex1
  1.2.3.4:34
end-set
!
prefix-set prefix_set_ex1
  10.0.0.0/16 ge 16 le 32,
  0.0.0.0/0 ge 25 le 32,
  0.0.0.0/0
end-set
!
route-policy policy_2
  if destination in prefix_set_ex1 then
    if (community matches-any com_set_ex1) then
      set community (10:666) additive
    endif
    if (extcommunity rt matches-any ext_comm_set_rt_ex1) then
      set community (10:999) additive
    endif
  endif
end-policy
!
```

Related Commands

Command	Description
show rpl maximum, on page 1839	Displays the maximum limits for lines of configuration and number of policies.

show rpl active as-path-set

To display the AS path sets that are referenced by at least one policy that is being used at an attach point, use the **show rpl active as-path-set** command in EXEC mode.

show rpl active as-path-set [**detail**]

Syntax Description

detail	(Optional) Displays the content of the object and all referenced objects for active AS path sets.

Command Default

No default behavior or values

Command Modes

EXEC

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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **show rpl active as-path-set** command to display all AS path sets that are in use in the system and referenced either directly or indirectly at a policy attach point.

Task ID

Task ID	Operations
route-policy	read

Examples

This example shows the following sample configuration:

```
router bgp 2
 address-family ipv4 unicast
 !
 neighbor 10.0.101.2
  remote-as 100
  address-family ipv4 unicast
   route-policy policy_1 in
  !
 !
 neighbor 10.0.101.3
  remote-as 12
  address-family ipv4 unicast
   route-policy policy_2 in
  !
 !
!
RP/0/RSP0/CPU0:router# show rpl route-policy policy_2 detail

prefix-set prefix_set_ex1
  10.0.0.0/16 ge 16 le 32,
  0.0.0.0/0 ge 25 le 32,
  0.0.0.0/0
```

```
end-set
!
community-set comm_set_ex1
  65500:1,
  65500:2,
  65500:3
end-set
!
extcommunity-set rt ext_comm_set_rt_ex1
   1.2.3.4:34
end-set
!
route-policy policy_2
   if destination in prefix_set_ex1 then
     if (community matches-any comm_set_ex1) then
       set community (10:666) additive
     endif
     if (extcommunity rt matches-any ext_comm_set_rt_ex1) then
       set community (10:999) additive
     endif
   endif
end-policy
!

RP/0/RSP0/CPU0:router# show rpl route-policy policy_1 detail

prefix-set prefix_set_ex1
  10.0.0.0/16 ge 16 le 32,
  0.0.0.0/0 ge 25 le 32,
  0.0.0.0/0
end-set
!
as-path-set as_path_set_ex1
  ios-regex '^_655--$',
  ios-regex '^_65501_$'
end-set
!
route-policy policy_1
  if (destination in prefix_set_ex1) then
    set local-preference 100
  endif
  if (as-path in as_path_set_ex1) then
    set community  (10:333) additive
  endif
end-policy
!
```

Given this sample configuration, the **show rpl active as-path-set** command displays the following information:

```
RP/0/RSP0/CPU0:router# show rpl active as-path-set

ACTIVE -- Referenced by at least one policy which is attached
INACTIVE -- Only referenced by policies which are not attached
UNUSED -- Not attached (directly or indirectly) and not referenced

The following as-path-sets are ACTIVE
-------------------------------------
as_path_set_ex1
```

Related Commands

Command	Description
show rpl active community-set, on page 1792	Displays the community sets that are referenced by at least one policy that is being used at an attach point.
show rpl active extcommunity-set, on page 1795	Displays the extended community sets that are referenced by at least one policy that is being used at an attach point.
show rpl active prefix-set, on page 1798	Displays the route policies that are referenced by at least one policy that is being used at an attach point.
show rpl active prefix-set, on page 1798	Displays the prefix sets that are referenced by at least one policy that is being used at an attach point.

show rpl active community-set

To display the community sets that are referenced by at least one policy that is being used at an attach point, use the **show rpl active community-set** command in EXEC mode.

show rpl active community-set [**detail**]

Syntax Description

detail	(Optional) Displays the content of the object and all referenced objects for active community sets.

Command Default

No default behavior or values

Command Modes

EXEC

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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **show rpl active community-set** command to display all community sets that are in use in the system and referenced either directly or indirectly at a policy attach point.

Task ID

Task ID	Operations
route-policy	read

Examples

This example shows the following sample configuration:

```
router bgp 2
 address-family ipv4 unicast
 !
 neighbor 10.0.101.2
  remote-as 100
  address-family ipv4 unicast
   route-policy policy_1 in
  !
 !
 neighbor 10.0.101.3
  remote-as 12
  address-family ipv4 unicast
   route-policy policy_2 in
  !
 !
!

RP/0/RSP0/CPU0:router# show rpl route-policy policy_2 detail

prefix-set prefix_set_ex1
  10.0.0.0/16 ge 16 le 32,
  0.0.0.0/0 ge 25 le 32,
```

```
  0.0.0.0/0
end-set
!
community-set comm_set_ex1
  65500:1,
  65500:2,
  65500:3
end-set
!
extcommunity-set rt ext_comm_set_rt_ex1
   1.2.3.4:34
end-set
!

route-policy policy_2
   if destination in prefix_set_ex1 then
     if (community matches-any comm_set_ex1) then
       set community (10:666) additive
     endif
     if (extcommunity rt matches-any ext_comm_set_rt_ex1) then
       set community (10:999) additive
     endif
   endif
end-policy
!

RP/0/RSP0/CPU0:router# show rpl route-policy policy_1 detail

prefix-set prefix_set_ex1
  10.0.0.0/16 ge 16 le 32,
  0.0.0.0/0 ge 25 le 32,
  0.0.0.0/0
end-set
!
as-path-set as_path_set_ex1
  ios-regex '^_655--$',
  ios-regex '^_65501_$'
end-set
!
route-policy policy_1
  if (destination in prefix_set_ex1) then
    set local-preference 100
  endif
  if (as-path in as_path_set_ex1) then
    set community  (10:333) additive
  endif
end-policy
!
```

Given this sample configuration, the **show rpl active community-set** command displays the following information:

```
RP/0/RSP0/CPU0:router# show rpl active community-set

ACTIVE -- Referenced by at least one policy which is attached
INACTIVE -- Only referenced by policies which are not attached
UNUSED -- Not attached (directly or indirectly) and not referenced

The following community-sets are ACTIVE
---------------------------------------
comm_set_ex1
```

Related Commands

Command	Description
show rpl active as-path-set, on page 1789	Displays the AS path sets that are referenced by at least one policy that is being used at an attach point.
show rpl active extcommunity-set, on page 1795	Displays the extended community sets that are referenced by at least one policy that is being used at an attach point.
show rpl active prefix-set, on page 1798	Displays the route policies that are referenced by at least one policy that is being used at an attach point.
show rpl active prefix-set, on page 1798	Displays the prefix sets that are referenced by at least one policy that is being used at an attach point.
show rpl active rd-set, on page 1801	Displays the route distinguisher sets that are referenced by at least one policy that is being used at an attach point.

show rpl active extcommunity-set

To display the extended community sets for cost, route target (RT), and Site-of-Origin (SoO) that are referenced by at least one route policy used at an attach point, use the **show rpl active extcommunity-set** command in EXEC mode.

show rpl active extcommunity-set [**cost** | **rt** | **soo**] [**detail**]

Syntax Description

cost	(Optional) Displays all extended community cost sets.
rt	(Optional) Displays all extended community RT sets.
soo	(Optional) Displays all extended community SoO sets.
detail	(Optional) Displays the content of the object and all referenced objects for active extended community sets.

Command Default

All extended community sets are displayed.

Command Modes

EXEC

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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **show rpl active extcommunity-set** command to display all extended community sets that are in use in the system and referenced either directly or indirectly at a policy attach point.

Task ID

Task ID	Operations
route-policy	read

Examples

This example shows the following sample configuration:

```
router bgp 2
 address-family ipv4 unicast
 !
 neighbor 10.0.101.2
  remote-as 100
  address-family ipv4 unicast
   route-policy policy_1 in
  !
 !
 neighbor 10.0.101.3
  remote-as 12
  address-family ipv4 unicast
```

```
   route-policy policy_2 in
  !
 !
!

RP/0/RSP0/CPU0:router# show rpl route-policy policy_2 detail

prefix-set prefix_set_ex1
  10.0.0.0/16 ge 16 le 32,
  0.0.0.0/0 ge 25 le 32,
  0.0.0.0/0
end-set
!
community-set comm_set_ex1
  65500:1,
  65500:2,
  65500:3
end-set
!
extcommunity-set rt ext_comm_set_rt_ex1
   1.2.3.4:34
end-set
!

route-policy policy_2
   if destination in prefix_set_ex1 then
     if (community matches-any comm_set_ex1) then
       set community (10:666) additive
     endif
     if (extcommunity rt matches-any ext_comm_set_rt_ex1) then
       set community (10:999) additive
     endif
   endif
end-policy
!

RP/0/RSP0/CPU0:router# show rpl route-policy policy_1 detail

prefix-set prefix_set_ex1
  10.0.0.0/16 ge 16 le 32,
  0.0.0.0/0 ge 25 le 32,
  0.0.0.0/0
end-set
!
as-path-set as_path_set_ex1
  ios-regex '^_655--$',
  ios-regex '^_65501_$'
end-set
!
route-policy policy_1
  if (destination in prefix_set_ex1) then
    set local-preference 100
  endif
  if (as-path in as_path_set_ex1) then
    set community  (10:333) additive
  endif
end-policy
!
```

Given this sample configuration, the **show rpl active extcommunity-set** command displays the following information:

```
RP/0/RSP0/CPU0:router# show rpl active extcommunity-set

ACTIVE -- Referenced by at least one policy which is attached
INACTIVE -- Only referenced by policies which are not attached

UNUSED -- Not attached (directly or indirectly) and not referenced

The following extcommunity-sets are ACTIVE
------------------------------------------
ext_comm_set_rt_ex1
```

Related Commands

Command	Description
show rpl active as-path-set, on page 1789	Displays the AS path sets that are referenced by at least one policy that is being used at an attach point.
show rpl active community-set, on page 1792	Displays the community sets that are referenced by at least one policy that is being used at an attach point.
show rpl active prefix-set, on page 1798	Displays the route policies that are referenced by at least one policy that is being used at an attach point.
show rpl active prefix-set, on page 1798	Displays the prefix sets that are referenced by at least one policy that is being used at an attach point.
show rpl active rd-set, on page 1801	Displays the route distinguisher sets that are referenced by at least one policy that is being used at an attach point.

show rpl active prefix-set

To display the prefix sets that are referenced by at least one policy that is being used at an attach point, use the **show rpl active prefix-set** command in EXEC mode.

show rpl active prefix-set [**detail**]

Syntax Description

detail	(Optional) Displays the content of the object and all referenced objects for active prefix sets.

Command Default

No default behavior or values

Command Modes

EXEC

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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **show rpl active prefix-set** command to display all prefix sets that are in use in the system and referenced either directly or indirectly at a policy attach point.

Task ID

Task ID	Operations
route-policy	read

Examples

This example shows the following sample configuration:

```
router bgp 2
 address-family ipv4 unicast
 !
 neighbor 10.0.101.2
  remote-as 100
  address-family ipv4 unicast
   route-policy policy_1 in
  !
 !
 neighbor 10.0.101.3
  remote-as 12
  address-family ipv4 unicast
   route-policy policy_2 in
  !
 !
!

RP/0/RSP0/CPU0:router# show rpl route-policy policy_2 detail

prefix-set prefix_set_ex1
  10.0.0.0/16 ge 16 le 32,
  0.0.0.0/0 ge 25 le 32,
```

```
  0.0.0.0/0
end-set
!

community-set comm_set_ex1
  65500:1,
  65500:2,
  65500:3
end-set
!
extcommunity-set rt ext_comm_set_rt_ex1
   1.2.3.4:34
end-set
!

route-policy policy_2
   if destination in prefix_set_ex1 then
     if (community matches-any comm_set_ex1) then
       set community (10:666) additive
     endif
     if (extcommunity rt matches-any ext_comm_set_rt_ex1) then
       set community (10:999) additive
     endif
   endif
end-policy
!

RP/0/RSP0/CPU0:router# show rpl route-policy policy_1 detail

prefix-set prefix_set_ex1
  10.0.0.0/16 ge 16 le 32,
  0.0.0.0/0 ge 25 le 32,
  0.0.0.0/0
end-set
!
as-path-set as_path_set_ex1
  ios-regex '^_655--$',
  ios-regex '^_65501_$'
end-set
!
route-policy policy_1
  if (destination in prefix_set_ex1) then
    set local-preference 100
  endif
  if (as-path in as_path_set_ex1) then
    set community  (10:333) additive
  endif
end-policy
!
```

The following example displays active prefix sets:

```
RP/0/RSP0/CPU0:router# show rpl active prefix-set

ACTIVE -- Referenced by at least one policy which is attached
INACTIVE -- Only referenced by policies which are not attached
UNUSED -- Not attached (directly or indirectly) and not referenced

The following prefix-sets are ACTIVE
------------------------------------
prefix_set_1
```

Related Commands

Command	Description
show rpl active as-path-set, on page 1789	Displays the AS path sets that are referenced by at least one policy that is being used at an attach point.
show rpl active community-set, on page 1792	Displays the community sets that are referenced by at least one policy that is being used at an attach point.
show rpl active extcommunity-set, on page 1795	Displays the extended community sets that are referenced by at least one policy that is being used at an attach point.
show rpl route-policy attachpoints, on page 1859	Displays the route policies that are referenced by at least one policy that is being used at an attach point.
show rpl active rd-set, on page 1801	Displays the route distinguisher sets that are referenced by at least one policy that is being used at an attach point.

show rpl active rd-set

To display the route distinguisher (RD) sets that are referenced by at least one policy that is being used at an attach point, use the **show rpl active rd-set** command in EXEC mode.

show rpl active rd-set [**detail**]

Syntax Description

detail	(Optional) Displays the content of the object and all referenced objects for active route policies.

Command Default

No default behavior or values

Command Modes

EXEC

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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **show rpl active rd-set** command to display all RD sets that are in use in the system and that are referenced either directly or indirectly at a policy attach point.

Task ID

Task ID	Operations
route-policy	read

Examples

This example shows the following sample configuration:

```
rd-set rdset1
   10:151,
   100.100.100.1:153,
   100.100.100.62/31:63
end-set
!
rd-set rdset2
   10:152,
   100.100.100.1:154,
   100.100.100.62/31:89
end-set
!
route-policy rdsetmatch
   if rd in rdset1 then
     set community (10:112)
   elseif rd in rdset2 then
     set community (10:223)
   endif
end-policy
!
router bgp 10
  bgp router-id 10.0.0.1
```

```
  address-family vpnv4 unicast
neighbor 10.10.10.1
   remote-as 10
   address-family ipv4 unicast
    route-policy rdsetmatch in
   !
  !
```

Given this sample configuration, the **show rpl active rd-set** command displays the following information:

```
RP/0/RSP0/CPU0:router# show rpl active rd-set

ACTIVE -- Referenced by at least one policy which is attached INACTIVE -- Only referenced
by policies which are not attached UNUSED -- Not attached (directly or indirectly) and not
 referenced

The following rd-sets are ACTIVE
----------------------------------------
     rdset1
     rdset2
```

Related Commands

Command	Description
show rpl active as-path-set, on page 1789	Displays the AS path sets that are referenced by at least one policy that is being used at an attach point.
show rpl active community-set, on page 1792	Displays the community sets that are referenced by at least one policy that is being used at an attach point.
show rpl active extcommunity-set, on page 1795	Displays the extended community sets that are referenced by at least one policy that is being used at an attach point.
show rpl active prefix-set, on page 1798	Displays the prefix sets that are referenced by at least one policy that is being used at an attach point.
show rpl active prefix-set, on page 1798	Displays the route policies that are referenced by at least one policy that is being used at an attach point.

show rpl active route-policy

To display the route policies that are referenced by at least one policy that is being used at an attach point, use the **show rpl active route-policy** command in EXEC mode.

show rpl active route-policy [**detail**]

Syntax Description

detail	(Optional) Displays the content of the object and all referenced objects for active route policies.

Command Default

No default behavior or values

Command Modes

EXEC

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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **show rpl active route-policy** command to display all policies that are in use in the system and that are referenced either directly or indirectly at a policy attach point.

Task ID

Task ID	Operations
route-policy	read

Examples

This example shows the following sample configuration:

```
router bgp 2
 address-family ipv4 unicast
 !
 neighbor 10.0.101.2
  remote-as 100
  address-family ipv4 unicast
   route-policy policy_1 in
  !
 !
 neighbor 10.0.101.3
  remote-as 12
  address-family ipv4 unicast
   route-policy policy_2 in
  !
 !
!

RP/0/RSP0/CPU0:router# show rpl route-policy policy_1

route-policy policy_1
  if (destination in prefix_set_ex1) then
    set local-preference 100
```

```
  endif
  if (as-path in as_path_set_ex1) then
    set community  (10:333) additive
  endif
end-policy
!
RP/0/RSP0/CPU0:router# show rpl route-policy policy_2

route-policy policy_2
   if destination in prefix_set_ex1 then
     if (community matches-any comm_set_ex1) then
       set community (10:666) additive
     endif
     if (extcommunity rt matches-any ext_comm_set_rt_ex1) then
       set community (10:999) additive
     endif
   endif
end-policy
!
```

Given this sample configuration, the **show rpl active route-policy** command displays the following information:

```
RP/0/RSP0/CPU0:router# show rpl active route-policy

ACTIVE -- Referenced by at least one policy which is attached
INACTIVE -- Only referenced by policies which are not attached
UNUSED -- Not attached (directly or indirectly) and not referenced

The following policies are (ACTIVE)
-----------------------------------
policy_1
policy_2
```

Related Commands

Command	Description
show rpl active as-path-set, on page 1789	Displays the AS path sets that are referenced by at least one policy that is being used at an attach point.
show rpl active community-set, on page 1792	Displays the community sets that are referenced by at least one policy that is being used at an attach point.
show rpl active extcommunity-set, on page 1795	Displays the extended community sets that are referenced by at least one policy that is being used at an attach point.
show rpl active prefix-set, on page 1798	Displays the prefix sets that are referenced by at least one policy that is being used at an attach point.
show rpl active rd-set, on page 1801	Displays the route distinguisher sets that are referenced by at least one policy that is being used at an attach point.

show rpl as-path-set

To display the contents of AS path sets, use the **show rpl as-path-set** command in EXEC mode.

show rpl as-path-set [*name* | **states** | **brief**]

Syntax Description

name	(Optional) Name of the AS path set.
states	(Optional) Displays all unused, inactive, and active states.
brief	(Optional) Limits the display to a list of the names of all AS path sets without their configurations.

Command Default

No default behavior or values

Command Modes

EXEC

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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the optional **brief** keyword to limit the display to a list of the names of all AS path sets without their configurations.

Task ID

Task ID	Operations
route-policy	read

Examples

This example shows the following sample configuration:

```
RP/0/RSP0/CPU0:router# show rpl route-policy policy_1

route-policy policy_1
  if (destination in prefix_set_ex1) then
    set local-preference 100
  endif
  if (as-path in as_path_set_ex1) then
    set community  (10:333) additive
  endif
end-policy
```

Given this sample configuration, the **show rpl as-path-set as_path_set_ex1** command displays the following information:

```
RP/0/RSP0/CPU0:router# show rpl as-path-set as_path_set_ex1

as-path-set as_path_set_ex1
```

```
  ios-regex '^_65500_$',
  ios-regex '^_65501_$'
end-set
```

Related Commands

Command	Description
show rpl community-set, on page 1812	Displays the configuration of a named community set.
show rpl extcommunity-set, on page 1819	Displays the configuration of a named extended community set.
show rpl route-policy, on page 1856	Displays the configuration of a named route policy.
show rpl prefix-set, on page 1843	Displays the configuration of a named prefix set.

show rpl as-path-set attachpoints

To display all of the policies used at an attach point that reference the named AS path set, use the **show rpl as-path-set attachpoints** command in EXEC mode.

show rpl as-path-set *name* **attachpoints**

Syntax Description

name	Name of an AS path set.

Command Default

No default behavior or values

Command Modes

EXEC

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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **show rpl as-path-set attachpoints** command to display all policies used at an attach point that reference the named set either directly or indirectly.

The AS path set name is required.

Task ID

Task ID	Operations
route-policy	read

Examples

This example shows the following sample configuration:

```
router bgp 2
 address-family ipv4 unicast
 !
 neighbor 10.0.101.2
  remote-as 100
  address-family ipv4 unicast
   route-policy policy_1 in
  !
 !
 neighbor 10.0.101.3
  remote-as 12
  address-family ipv4 unicast
   route-policy policy_2 in
  !
 !
!

RP/0/RSP0/CPU0:router# show rpl route-policy policy_1

route-policy policy_1
```

```
  if (destination in prefix_set_ex1) then
    set local-preference 100
  endif
  if (as-path in as_path_set_ex1) then
    set community  (10:333) additive
  endif
end-policy
!
RP/0/RSP0/CPU0:router# show rpl route-policy policy_2

route-policy policy_2
  if (destination in prefix_set_ex1) then
    if (community matches-any comm_set_ex1) then
      set community  (10:666) additive
    endif
    if (extcommunity matches-any ext_comm_set_rt_ex1) then
      set community  (10:999) additive
    endif
  endif
end-policy
!
```

Given this sample configuration, the **show rpl as-path-set as_path_set_ex1 attachpoints** command displays the following information:

```
RP/0/RSP0/CPU0:router# show rpl as-path-set as_path_set_ex1 attachpoints

BGP Attachpoint:Neighbor

Neighbor/Group  type  afi/safi   in/out     referring policy attached policy
-------------------------------------------------------------------------------
10.0.101.2      --    IPv4/uni    in         policy_1         policy_1
10.0.101.3      --    IPv4/uni    in         policy_2         policy_2
```

This table describes the significant fields shown in the display.

Table 175: show rpl as-path-set attachpoints Field Descriptions

Field	Description
BGP Attachpoint	Location of the attach point.
Neighbor/Group	IP address of the attach point on the neighbor.
type	Displays the address family mode.
afi/safi	Address family identifier or subsequent address family identifier.
in/out	Import or export policy.
referring policy	Policy that refers to the AS path set.
attached policy	Policy used at the attach point.

Related Commands

Command	Description
show rpl community-set attachpoints, on page 1814	Displays all the policies used at an attach point that reference the named community set.
show rpl route-policy attachpoints, on page 1859	Displays all the policies used at an attach point that reference the named policy.
show rpl prefix-set attachpoints, on page 1845	Displays all the policies used at an attach point that reference the named prefix set.

show rpl as-path-set references

To list all of the policies that reference the named AS path set, use the **show rpl as-path-set references** command in EXEC mode.

show rpl as-path-set *name* **references** [**brief**]

Syntax Description

name	Name of the prefix set.
brief	(Optional) Limits the output to just the brief table and not the detailed information for the named AS path set.

Command Default

No default behavior or values

Command Modes

EXEC

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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **show rpl as-path-set references** command to display all policies that reference the named AS path set either directly or indirectly.

Use the optional **brief** keyword to limit the output to just a summary table and not the detailed information for the AS path set.

Task ID

Task ID	Operations
route-policy	read

Examples

This example shows the following sample configuration:

```
router bgp 2
 address-family ipv4 unicast
 !
 neighbor 10.0.101.2
  remote-as 100
  address-family ipv4 unicast
   route-policy policy_1 in
  !
 !
RP/0/RSP0/CPU0:router# show rpl route-policy policy_1

route-policy policy_1
  if (destination in prefix_set_ex1) then
    set local-preference 100
  endif
```

```
  if (as-path in as_path_set_ex1) then
    set community  (10:333) additive
  endif
end-policy
```

Given this sample configuration, the **show rpl as-path-set as_path_set_ex1 references** command displays the following information:

```
RP/0/RSP0/CPU0:router# show rpl as-path-set as_path_set_ex1 references

Usage Direct -- Reference occurs in this policy
Usage Indirect -- Reference occurs via an apply statement

Status UNUSED -- Policy is not in use at an attachpoint (unattached)
Status ACTIVE -- Policy is actively used at an attachpoint
Status INACTIVE -- Policy is applied by an unattached policy

     Usage/Status         count
-----------------------------------------------------------------
     Direct               1
     Indirect             0

     ACTIVE               1
     INACTIVE             0
     UNUSED               0

     route-policy         usage      policy status
-----------------------------------------------------------------
    policy_1             Direct     ACTIVE
```

This table describes the significant fields shown in the display.

Table 176: show rpl as-path-set references Field Descriptions

Field	Description
Usage/Status	Displays the usage and status of all policies that reference the AS path set. Values for usage are Direct or Indirect. Values for policy status are ACTIVE, INACTIVE, or UNUSED.
count	Number of policies that match each usage and status option.
route-policy	Name of the route policies that reference the AS path set.
usage	Type of usage for the policy.
policy status	Status of the policy.

Related Commands

Command	Description
show rpl community-set references, on page 1816	Lists all policies that reference the named community set.
show rpl route-policy references, on page 1864	Lists all policies that reference the named policy.
show rpl prefix-set references, on page 1848	Lists all policies that reference the named prefix set.

show rpl community-set

To display the configuration of community sets, use the **show rpl community-set** command in EXEC mode.

show rpl community-set [*name* | **states** | **brief**]

Syntax Description

name	(Optional) Name of the community set.
states	(Optional) Shows all unused, inactive, and active states.
brief	(Optional) Limits the display to a list of the names of all community sets without their configurations.

Command Default

No default behavior or values

Command Modes

EXEC

Command History

Release	Modification
Release 3.7.2	This command was introduced.
Release 5.3.2	The command output was modified to display graceful maintenance feature information.

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.

Use the optional **brief** keyword to limit the display to a list of the names of community sets without their configurations.

Task ID

Task ID	Operations
route-policy	read

The following is the sample output of the show rpl community-set command with graceful maintenance feature attributes displayed:

```
RP/0/0/CPU0:R5#show rpl community-set
Thu Jan 29 17:55:04.792 PST
Listing for all Community Set objects

community-set gshut
  graceful-shutdown
end-set
```

Examples

This example shows the following sample configuration:

```
route-policy policy_4
  if (destination in prefix_set_ex2) then
    if (community matches-any comm_set_ex2) then
      set community  (10:666) additive
```

```
    endif
    if (extcommunity matches-any ext_comm_set_rt_ex2) then
      set community  (10:999) additive
    endif
  endif
end-policy
```

Given this sample configuration, the **show rpl community-set comm_set_ex2** command displays the following information:

```
RP/0/RSP0/CPU0:router# show rpl community-set comm_set_ex2

community-set comm_set_ex2
  65501:1,
  65501:2,
  65501:3
end-set
```

Related Commands

Command	Description
show rpl as-path-set, on page 1805	Displays the configuration of a named AS path set.
show rpl extcommunity-set, on page 1819	Displays the configuration of a named extended community set.
show rpl prefix-set, on page 1843	Displays the configuration of a named prefix set.
show rpl rd-set, on page 1850	Displays the configuration of a named RD set.
show rpl route-policy, on page 1856	Displays the configuration of a named route policy.

show rpl community-set attachpoints

To display all the policies used at an attach point that reference the named community set, use the **show rpl community-set attachpoints** command in EXEC mode.

show rpl community-set *name* **attachpoints**

Syntax Description

name	Name of a community set.

Command Default

No default behavior or values

Command Modes

EXEC

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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **show rpl community-set attachpoints** command to display all the policies used at an attach point that reference the named community set either directly or indirectly.

The community set name is required.

Task ID

Task ID	Operations
route-policy	read

Examples

This example shows the following sample configuration:

```
router bgp 2
 address-family ipv4 unicast
 !
 neighbor 10.0.101.3
  remote-as 12
  address-family ipv4 unicast
   route-policy policy_2 in
  !
 !
!
!
route-policy policy_2
   if destination in prefix_set_ex1 then
     if (community matches-any comm_set_ex1) then
       set community (10:666) additive
     endif
     if (extcommunity rt matches-any ext_comm_set_rt_ex1) then <<<<<
       set community (10:999) additive
     endif
   endif
end-policy
```

```
!
```

Given this sample configuration, the **show rpl community-set attachpoints** command displays the following information:

```
RP/0/RSP0/CPU0:router# show rpl community-set ext_comm_set_rt_ex1 attachpoints

BGP Attachpoint:Neighbor

Neighbor/Group  type  afi/safi   in/out     referring policy attached policy
-------------------------------------------------------------------------------
10.0.101.3      --    IPv4/uni   in         policy_2         policy_2
```

This table describes the significant fields shown in the display.

Table 177: show rpl community-set attachpoints Field Descriptions

Field	Description
BGP Attachpoint	Location of the attach point.
Neighbor/Group	IP address of the attach point on the neighbor.
type	Displays the address family mode.
afi/safi	Address family identifier or subsequent address family identifier.
in/out	Import or export policy.
referring policy	Policy that refers to the AS path set.
attached policy	Policy used at the attach point.

Related Commands

Command	Description
show rpl as-path-set attachpoints, on page 1807	Displays all the policies used at an attach point that reference the named AS path set.
show rpl prefix-set attachpoints, on page 1845	Displays all the policies used at an attach point that reference the named prefix set.
show rpl rd-set attachpoints, on page 1852	Displays all the policies used at an attach point that reference the named RD set.
show rpl route-policy attachpoints, on page 1859	Displays all the policies used at an attach point that reference the named policy.

show rpl community-set references

To list all the policies that reference the named community set, use the **show rpl community-set references** command in EXEC mode.

show rpl community-set *name* **references** [**brief**]

Syntax Description

name	Name of a community set.
brief	(Optional) Limits the output to just the summary table and not the detailed information for the community set.

Command Default

No default behavior or values

Command Modes

EXEC

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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **show rpl community-set references** command to display all the policies that reference the named community set.

Use the optional **brief** keyword to limit the output to just a summary table and not the detailed information for the community set.

Task ID

Task ID	Operations
route-policy	read

Examples

This example shows the following sample configuration:

```
router bgp 2
 address-family ipv4 unicast
 !
 neighbor 10.0.101.3
  remote-as 12
  address-family ipv4 unicast
   route-policy policy_2 in
  !
 !
!
route-policy policy_2
  if (destination in prefix_set_ex1) then
    if (community matches-any comm_set_ex1) then
      set community  (10:666) additive
    endif
```

```
    if (extcommunity matches-any ext_comm_set_rt_ex1) then
      set community  (10:999) additive
    endif
  endif
end-policy
```

Given this sample configuration, the **show rpl extcommunity-set comm_set_ex1 references** command displays the following information:

```
RP/0/RSP0/CPU0:router# show rpl extcommunity-set comm_set_ex1 references

Usage Direct -- Reference occurs in this policy
Usage Indirect -- Reference occurs via an apply statement

Status UNUSED -- Policy is not in use at an attachpoint (unattached)
Status ACTIVE -- Policy is actively used at an attachpoint
Status INACTIVE -- Policy is applied by an unattached policy

     Usage/Status         count
-------------------------------------------------------------
     Direct               1
     Indirect             0

     ACTIVE               1
     INACTIVE             0
     UNUSED               0


     route-policy         usage      policy status
-------------------------------------------------------------
    policy_2             Direct     ACTIVE
```

This table describes the significant fields shown in the display.

Table 178: show rpl community-set references Field Descriptions

Field	Description
Usage/Status	Displays the usage and status of all policies that reference the community set. Values for usage are Direct or Indirect. Values for status are ACTIVE, INACTIVE, and UNUSED.
count	Number of policies that match each usage and status option.
route-policy	Name of the route policies that reference the community set.
usage	Type of usage for the policy.
policy status	Status of the policy.

Related Commands

Command	Description
show rpl as-path-set references, on page 1810	Lists all policies that reference the named AS path set.
show rpl prefix-set references, on page 1848	Lists all policies that reference the named prefix set.

Command	Description
show rpl rd-set references, on page 1854	Lists all policies that reference the named RD set.
show rpl route-policy references, on page 1864	Lists all policies that reference the named policy.

show rpl extcommunity-set

To display the configuration of extended community sets, use the **show rpl extcommunity-set** command in EXEC mode.

show rpl extcommunity-set [*name* [**attachpoints** | **references**]] [**cost** | **rt** | **soo**] [*name*] [**brief**] [**states**]

Syntax Description

name	(Optional) Name of the community set.
attachpoints	(Optional) Displays all attach points for this community set.
references	(Optional) Displays all policies that use this community set.
cost	(Optional) Displays all extended community cost sets.
rt	(Optional) Displays all extended community RT sets.
soo	(Optional) Displays all extended community SoO sets.
brief	(Optional) Limits the display to a list of the names of all extended community sets without their configurations.
states	(Optional) Displays all unused, inactive, and active states.

Command Default

If an attachpoint or reference is not specified, all configured extended community sets are displayed

If a cost, RT, or SoO sets is not specified, all configured extended community sets are displayed

Command Modes

EXEC

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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the optional **brief** keyword to limit the display to a list of the names of extended community sets without their configurations.

Task ID

Task ID	Operations
route-policy	read

Examples

In the following example, the configuration of an extended community is displayed for the RT community set named ext_comm_set_rt_ex1:

```
RP/0/RSP0/CPU0:router# show rpl extcommunity-set rt ext_comm_set_rt_ex1
```

```
ext_comm_set_rt_ex1
   1.2.3.4:34
end-set
!
```

In the following example, the configuration of an extended community is displayed with all RT set objects:

```
RP/0/RSP0/CPU0:router# show rpl extcommunity-set rt

Listing for all Extended Community RT Set objects

extcommunity-set rt extrt1
  66:60001
end-set
!
extcommunity-set rt rtset1
  10:615,
  10:6150,
  15.15.15.15:15
end-set
!
extcommunity-set rt rtset3
  11:11,
  11.1.1.1:3
end-set
!
extcommunity-set rt extsoo1
  66:70001
end-set
!
extcommunity-set rt rtset11
  100:121,
  100:122,
  100:123,
  100:124,
  100:125,
  100:126,
  100:127,
  100:128,
  7.7.7.7:21
end-set
!
```

In the following example, the configuration of an extended community is displayed with all cost set objects:

```
RP/0/RSP0/CPU0:router# show rpl extcommunity-set cost

Listing for all Extended Community COST Set objects

extcommunity-set cost costset1
  IGP:90:914,
  Pre-Bestpath:91:915
end-set
!
extcommunity-set cost costset2
  IGP:92:916,
  Pre-Bestpath:93:917,
  IGP:94:918,
  Pre-Bestpath:95:919
```

```
end-set
!
```

In the following example, the configuration of an extended community is displayed with all SoO set objects:

```
Extended Community SOO Set objects

extcommunity-set soo sooset1
  10:151,
  100.100.100.1:153
end-set
!
extcommunity-set soo sooset3
  11:11,
  11.1.1.1:3
end-set
!
```

Related Commands

Command	Description
show rpl as-path-set, on page 1805	Displays the configuration of a named AS path set.
show rpl community-set, on page 1812	Displays the configuration of a named community set.
show rpl prefix-set, on page 1843	Displays the configuration of a named prefix set.
show rpl rd-set, on page 1850	Displays the configuration of a named RD set.
show rpl route-policy, on page 1856	Displays the configuration of a named route policy.

show rpl inactive as-path-set

To display the AS path sets that are referenced by a policy but not in any policy that is used at an attach point, use the **show rpl inactive as-path-set** command in EXEC mode.

show rpl inactive as-path-set [**detail**]

Syntax Description

detail	(Optional) Displays the content of the object and all referenced objects for inactive AS path sets.

Command Default

No default behavior or values

Command Modes

EXEC

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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **show rpl inactive as-path-set** command to display all AS path sets that are not in use at an attach point either directly or indirectly but are referenced by at least one policy in the system.

Task ID

Task ID	Operations
route-policy	read

Examples

This example shows the following sample configuration:

```
router bgp 2
 address-family ipv4 unicast
 !
 neighbor 10.0.101.2
  remote-as 100
  address-family ipv4 unicast
   route-policy policy_1 in
  !
 !
 neighbor 10.0.101.3
  remote-as 12
  address-family ipv4 unicast
   route-policy policy_2 in
  !
 !
!
route-policy sample
  if (destination in sample) then
    drop
  endif
end-policy
!
```

```
route-policy policy_1
  if (destination in prefix_set_ex1) then
    set local-preference 100
  endif
  if (as-path in as_path_set_ex1) then
    set community  (10:333) additive
  endif
end-policy
!
route-policy policy_2
   if destination in prefix_set_ex1 then
     if (community matches-any comm_set_ex1) then
       set community (10:666) additive
     endif
     if (extcommunity rt matches-any ext_comm_set_rt_ex1) then
       set community (10:999) additive
     endif
   endif
end-policy
!
route-policy policy_3
  if (destination in prefix_set_ex2) then
    set local-preference 100
  endif
  if (as-path in as_path_set_ex2) then
    set community  (10:333) additive
  endif
end-policy
!
route-policy policy_4
  if (destination in prefix_set_ex2) then
    if (community matches-any comm_set_ex2) then
      set community  (10:666) additive
    endif
    if (extcommunity matches-any ext_comm_set_rt_ex2) then
      set community  (10:999) additive
    endif
  endif
end-policy
!
route-policy policy_5
  apply sample1
  apply policy_3
end-policy
```

Given this sample configuration, the **show rpl inactive as-path-set** command displays the following information:

```
RP/0/RSP0/CPU0:router# show rpl inactive as-path-set

ACTIVE -- Referenced by at least one policy which is attached
INACTIVE -- Only referenced by policies which are not attached
UNUSED -- Not attached (directly or indirectly) and not referenced

The following as-path-sets are INACTIVE
---------------------------------------
as_path_set_ex2
```

Related Commands

Command	Description
show rpl inactive community-set, on page 1825	Displays the community sets that are referenced by a policy but not in any policy that is used at an attach point.
show rpl inactive extcommunity-set, on page 1828	Displays the extended community sets that are referenced by a policy but not in any policy that is used at an attach point.
show rpl inactive prefix-set, on page 1831	Displays the prefix sets that are referenced by a policy but not in any policy that is used at an attach point.
show rpl inactive rd-set, on page 1834	Displays the RD sets that are referenced by a policy but not in any policy that is used at an attach point.
show rpl inactive route-policy, on page 1836	Displays the route policies that are referenced by a policy but not in any policy that is used at an attach point.

show rpl inactive community-set

To display the community sets that are referenced by a policy but not any policy that is used at an attach point, use the **show rpl inactive community-set** command in EXEC mode.

show rpl inactive community-set [**detail**]

Syntax Description

detail	(Optional) Displays the content of the object and all referenced objects for inactive community sets.

Command Default

No default behavior or values

Command Modes

EXEC

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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **show rpl inactive community-set** command to display all community sets that are not in use at an attach point either directly or indirectly but are referenced by at least one policy in the system.

Task ID

Task ID	Operations
route-policy	read

Examples

This example shows the following sample configuration:

```
router bgp 2
 address-family ipv4 unicast
 !
 neighbor 10.0.101.2
  remote-as 100
  address-family ipv4 unicast
   route-policy policy_1 in
  !
 !
 neighbor 10.0.101.3
  remote-as 12
  address-family ipv4 unicast
   route-policy policy_2 in
  !
 !
!
route-policy sample2
  if (destination in sample2) then
    drop
  endif
end-policy
!
```

```
route-policy policy_1
  if (destination in prefix_set_ex1) then
    set local-preference 100
  endif
  if (as-path in as_path_set_ex1) then
    set community  (10:333) additive
  endif
end-policy
!
route-policy policy_2
   if destination in prefix_set_ex1 then
     if (community matches-any comm_set_ex1) then
       set community (10:666) additive
     endif
     if (extcommunity rt matches-any ext_comm_set_rt_ex1) then
       set community (10:999) additive
     endif
   endif
end-policy
!
route-policy policy_3
  if (destination in prefix_set_ex2) then
    set local-preference 100
  endif
  if (as-path in as_path_set_ex2) then
    set community  (10:333) additive
  endif
end-policy
!
route-policy policy_4
  if (destination in prefix_set_ex2) then
    if (community matches-any comm_set_ex2) then
      set community  (10:666) additive
    endif
    if (extcommunity matches-any ext_comm_set_rt_ex2) then
      set community  (10:999) additive
    endif
  endif
end-policy
!
route-policy policy_5
  apply sample2
  apply policy_3
end-policy
```

Given this sample configuration, the **show rpl inactive community-set** command displays the following information:

```
RP/0/RSP0/CPU0:router# show rpl inactive community-set

ACTIVE -- Referenced by at least one policy which is attached
INACTIVE -- Only referenced by policies which are not attached
UNUSED -- Not attached (directly or indirectly) and not referenced

The following community-sets are INACTIVE
-----------------------------------------
comm_set_ex2
```

Related Commands

Command	Description
show rpl inactive as-path-set, on page 1822	Displays the AS path sets that are referenced by a policy but not in any policy that is used at an attach point.
show rpl inactive extcommunity-set, on page 1828	Displays the extended community sets that are referenced by a policy but not in any policy that is used at an attach point.
show rpl inactive prefix-set, on page 1831	Displays the prefix sets that are referenced by a policy but not in any policy that is used at an attach point.
show rpl inactive rd-set, on page 1834	Displays the RD sets that are referenced by a policy but not in any policy that is used at an attach point.
show rpl inactive route-policy, on page 1836	Displays the route policies that are referenced by a policy but not in any policy that is used at an attach point.

show rpl inactive extcommunity-set

To display the extended community sets that are referenced by a policy but not in any policy that is used at an attach point, use the **show rpl inactive extcommunity-set** command in EXEC mode.

show rpl inactive extcommunity-set [**detail**]

Syntax Description

detail	(Optional) Displays the content of the object and all referenced objects for inactive extended community sets.

Command Default

No default behavior or values

Command Modes

EXEC

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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **show rpl inactive extcommunity-set** command to display all extended community sets that are not in use at an attach point either directly or indirectly but are referenced by at least one policy in the system.

Task ID

Task ID	Operations
route-policy	read

Examples

This example shows the following sample configuration:

```
router bgp 2
 address-family ipv4 unicast
 !
 neighbor 10.0.101.2
  remote-as 100
  address-family ipv4 unicast
   route-policy policy_1 in
  !
 !
 neighbor 10.0.101.3
  remote-as 12
  address-family ipv4 unicast
   route-policy policy_2 in
  !
 !
!
route-policy sample3
  if (destination in sample3) then
    drop
  endif
```

```
end-policy
!
route-policy policy_1
  if (destination in prefix_set_ex1) then
    set local-preference 100
  endif
  if (as-path in as_path_set_ex1) then
    set community  (10:333) additive
  endif
end-policy
!
route-policy policy_2
   if destination in prefix_set_ex1 then
     if (community matches-any comm_set_ex1) then
       set community (10:666) additive
     endif
     if (extcommunity rt matches-any ext_comm_set_rt_ex1) then
       set community (10:999) additive
     endif
   endif
end-policy
!
route-policy policy_3
  if (destination in prefix_set_ex2) then
    set local-preference 100
  endif
  if (as-path in as_path_set_ex2) then
    set community  (10:333) additive
  endif
end-policy
!
route-policy policy_4
  if (destination in prefix_set_ex2) then
    if (community matches-any comm_set_ex2) then
      set community  (10:666) additive
    endif
    if (extcommunity matches-any ext_comm_set_rt_ex2) then
      set community  (10:999) additive
    endif
  endif
end-policy
!
route-policy policy_5
  apply sample3
  apply policy_3
end-policy
```

Given this sample configuration, the **show rpl inactive extcommunity-set** command displays the following information:

```
RP/0/RSP0/CPU0:router# show rpl inactive extcommunity-set

ACTIVE -- Referenced by at least one policy which is attached
INACTIVE -- Only referenced by policies which are not attached
UNUSED -- Not attached (directly or indirectly) and not referenced

The following extcommunity-sets are INACTIVE
-------------------------------------------
ext_comm_set_rt_ex2
```

Related Commands

Command	Description
show rpl inactive as-path-set, on page 1822	Displays the AS path sets that are referenced by a policy but not in any policy that is used at an attach point.
show rpl inactive community-set, on page 1825	Displays the community sets that are referenced by a policy but not in any policy that is used at an attach point.
show rpl inactive prefix-set, on page 1831	Displays the prefix sets that are referenced by a policy but not in any policy that is used at an attach point.
show rpl inactive rd-set, on page 1834	Displays the RD sets that are referenced by a policy but not in any policy that is used at an attach point.
show rpl inactive route-policy, on page 1836	Displays the route policies that are referenced by a policy but not in any policy that is used at an attach point.

show rpl inactive prefix-set

To display the prefix sets that are referenced by a policy but not in any policy that is used at an attach point, use the **show rpl inactive prefix-set** command in EXEC mode.

show rpl inactive prefix-set [**detail**]

Syntax Description

detail	(Optional) Displays the content of the object and all referenced objects for inactive prefix sets.

Command Default

No default behavior or values

Command Modes

EXEC

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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **show rpl inactive prefix-set** command to display all prefix sets that are not in use at an attach point either directly or indirectly but are referenced by at least one policy in the system.

Task ID

Task ID	Operations
route-policy	read

Examples

This example shows the following sample configuration:

```
router bgp 2
 address-family ipv4 unicast
 !
 neighbor 10.0.101.2
  remote-as 100
  address-family ipv4 unicast
   route-policy policy_1 in
  !
 !
 neighbor 10.0.101.3
  remote-as 12
  address-family ipv4 unicast
   route-policy policy_2 in
  !
 !
!
route-policy sample4
  if (destination in sample4) then
    drop
  endif
end-policy
!
```

```
route-policy policy_1
  if (destination in prefix_set_ex1) then
    set local-preference 100
  endif
  if (as-path in as_path_set_ex1) then
    set community  (10:333) additive
  endif
end-policy
!
route-policy policy_2
   if destination in prefix_set_ex1 then
     if (community matches-any comm_set_ex1) then
       set community (10:666) additive
     endif
     if (extcommunity rt matches-any ext_comm_set_rt_ex1) then
       set community (10:999) additive
     endif
   endif
end-policy
!
route-policy policy_3
  if (destination in prefix_set_ex2) then
    set local-preference 100
  endif
  if (as-path in as_path_set_ex2) then
    set community  (10:333) additive
  endif
end-policy
!
route-policy policy_4
  if (destination in prefix_set_ex2) then
    if (community matches-any comm_set_ex2) then
      set community  (10:666) additive
    endif
    if (extcommunity matches-any ext_comm_set_rt_ex2) then
      set community  (10:999) additive
    endif
  endif
end-policy
!
route-policy policy_5
  apply sample4
  apply policy_3
end-policy
```

Given this sample configuration, the **show rpl inactive prefix-set** command displays the following information:

```
RP/0/RSP0/CPU0:router# show rpl inactive prefix-set

ACTIVE -- Referenced by at least one policy which is attached
INACTIVE -- Only referenced by policies which are not attached
UNUSED -- Not attached (directly or indirectly) and not referenced

The following prefix-sets are INACTIVE
--------------------------------------
sample4
prefix_set_ex2
```

Related Commands

Command	Description
show rpl inactive as-path-set, on page 1822	Displays the AS path sets that are referenced by a policy but not in any policy that is used at an attach point.
show rpl inactive community-set, on page 1825	Displays the community sets that are referenced by a policy but not in any policy that is used at an attach point.
show rpl inactive extcommunity-set, on page 1828	Displays the extended community sets that are referenced by a policy but not in any policy that is used at an attach point.
show rpl inactive route-policy, on page 1836	Displays the route policies that are referenced by a policy but not in any policy that is used at an attach point.
show rpl inactive rd-set, on page 1834	Displays the RD sets that are referenced by a policy but not in any policy that is used at an attach point.

show rpl inactive rd-set

To display the route distinguisher (RD) sets that are referenced by a policy but not in any policy that is used at an attach point, use the **show rpl inactive rd-set** command in EXEC mode.

show rpl inactive rd-set [**detail**]

Syntax Description

detail	(Optional) Displays the content of the object and all referenced objects for inactive RD sets.

Command Default

No default behavior or values

Command Modes

EXEC

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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **show rpl inactive rd-set** command to display all RD sets that are not in use at an attach point either directly or indirectly but are referenced by at least one policy in the system.

Task ID

Task ID	Operations
route-policy	read

Examples

This example shows the following sample configuration:

```
rd-set rdset1
  10:151,
  100.100.100.1:153,
  100.100.100.62/31:63
end-set
!
rd-set rdset2
  10:152,
  100.100.100.1:154,
  100.100.100.62/31:89
end-set
!
```

Given this sample configuration, the **show rpl inactive rd-set** command displays the following information:

```
RP/0/RSP0/CPU0:router# show rpl inactive rd-set

ACTIVE -- Referenced by at least one policy which is attached INACTIVE -- Only referenced
by policies which are not attached UNUSED -- Not attached (directly or indirectly) and not
 referenced
```

```
The following rd-sets are INACTIVE
----------------------------------------
    rdset1
    rdset2
```

Related Commands

Command	Description
show rpl inactive as-path-set, on page 1822	Displays the AS path sets that are referenced by a policy but not in any policy that is used at an attach point.
show rpl inactive community-set, on page 1825	Displays the community sets that are referenced by a policy but not in any policy that is used at an attach point.
show rpl inactive extcommunity-set, on page 1828	Displays the extended community sets that are referenced by a policy but not in any policy that is used at an attach point.
show rpl inactive prefix-set, on page 1831	Displays the prefix sets that are referenced by a policy but not in any policy that is used at an attach point.
show rpl inactive route-policy, on page 1836	Displays the route policies that are referenced by a policy but not in any policy that is used at an attach point.

show rpl inactive route-policy

To display the route policies that are referenced by a policy but not in any policy that is used at an attach point, use the **show rpl inactive route-policy** command in EXEC mode.

show rpl inactive route-policy [**detail**]

Syntax Description

detail	(Optional) Displays the content of the object and all referenced objects for inactive route policies.

Command Default

No default behavior or values

Command Modes

EXEC

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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **show rpl inactive route-policy** command to display all policies that are not in use at an attach point either directly or indirectly but are referenced by at least one other policy in the system.

Task ID

Task ID	Operations
route-policy	read

Examples

This example shows the following sample configuration:

```
router bgp 2
 address-family ipv4 unicast
 !
 neighbor 10.0.101.2
  remote-as 100
  address-family ipv4 unicast
   route-policy policy_1 in
  !
 !
 neighbor 10.0.101.3
  remote-as 12
  address-family ipv4 unicast
   route-policy policy_2 in
  !
 !
!
route-policy sample3
  if (destination in sample3) then
    drop
  endif
end-policy
!
```

```
route-policy policy_1
  if (destination in prefix_set_ex1) then
    set local-preference 100
  endif
  if (as-path in as_path_set_ex1) then
    set community  (10:333) additive
  endif
end-policy
!
route-policy policy_2
   if destination in prefix_set_ex1 then
     if (community matches-any comm_set_ex1) then
       set community (10:666) additive
     endif
     if (extcommunity rt matches-any ext_comm_set_rt_ex1) then
       set community (10:999) additive
     endif
   endif
end-policy
!
route-policy policy_3
  if (destination in prefix_set_ex2) then
    set local-preference 100
  endif
  if (as-path in as_path_set_ex2) then
    set community  (10:333) additive
  endif
end-policy
!
route-policy policy_4
  if (destination in prefix_set_ex2) then
    if (community matches-any comm_set_ex2) then
      set community  (10:666) additive
    endif
    if (extcommunity matches-any ext_comm_set_rt_ex2) then
      set community  (10:999) additive
    endif
  endif
end-policy
!
route-policy policy_5
  apply sample3
  apply policy_3
end-policy
```

Given this sample configuration, the **show rpl inactive route-policy** command displays the following information:

```
RP/0/RSP0/CPU0:router# show rpl inactive route-policy

ACTIVE -- Referenced by at least one policy which is attached
INACTIVE -- Only referenced by policies which are not attached
UNUSED -- Not attached (directly or indirectly) and not referenced

The following policies are (INACTIVE)
-------------------------------------
sample3
policy_3
```

Related Commands

Command	Description
show rpl inactive as-path-set, on page 1822	Displays the AS path sets that are referenced by a policy but not in any policy that is used at an attach point.
show rpl inactive community-set, on page 1825	Displays the community sets that are referenced by a policy but not in any policy that is used at an attach point.
show rpl inactive extcommunity-set, on page 1828	Displays the extended community sets that are referenced by a policy but not in any policy that is used at an attach point.
show rpl inactive prefix-set, on page 1831	Displays the prefix sets that are referenced by a policy but not in any policy that is used at an attach point.
show rpl inactive rd-set, on page 1834	Displays the RD sets that are referenced by a policy but not in any policy that is used at an attach point.

show rpl maximum

To display the maximum limits for lines of configuration and number of policies, use the **show rpl maximum** command in EXEC mode.

show rpl maximum [**lines** | **policies**]

Syntax Description

lines	(Optional) Displays the number of lines of configuration limit.
policies	(Optional) Displays the number of policies limit.

Command Default

No default behavior or values

Command Modes

EXEC

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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **show rpl maximum** command to display the current total, current limit, and maximum limit for lines of configuration and policies.

Use the optional **lines** keyword to limit the display to the number of lines of configuration limits. Use the optional **policies** keyword to limit the display to the number of policies limits.

Task ID

Task ID	Operations
route-policy	read

Examples

The following example shows sample output from the **show rpl maximum** command:

```
RP/0/RSP0/CPU0:router# show rpl maximum
                              Current    Current       Max
                               Total      Limit       Limit
-------------------------------------------------------------
Lines of configuration             3      65536     131072
Policies                           1       3500       5000
Compiled policies size (kB)        0
```

Table 179: show rpl maximum Field Descriptions, on page 1840 describes the significant fields shown in the display.

Table 179: show rpl maximum Field Descriptions

Field	Description
Lines of configuration	Displays the current total, current limit, and maximum limit of lines for the policy.
Policies	Displays the current total, current limit, and maximum limit of policies.
Compiled policies size (kB)	Displays the current compiled total for policies in kilobytes.

Related Commands

Command	Description
rpl maximum, on page 1735	Configures the maximum number of lines of configuration and number of policies.

show rpl policy-global references

To display policy-global definitions, use the **show rpl policy-global references** command in EXEC mode.

show rpl policy-global references [**brief**]

Syntax Description

brief	(Optional) Limits the display to a list of the policy names.

Command Default

No default behavior or values

Command Modes

EXEC

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 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
route-policy	read

Examples

This example shows the following sample configuration:

```
policy-global
   infinity '16'
end-global
!
route-policy set-rip-unreachable
   set rip-metric $infinity
end-policy
!
```

Given this sample configuration, the **show rpl policy-global references** command displays the following information:

```
RP/0/RSP0/CPU0:router# show rpl policy-global references

Usage Direct -- Reference occurs in this policy Usage Indirect -- Reference occurs via an
apply statement

Status UNUSED -- Policy is not in use at an attachpoint (unattached) Status ACTIVE -- Policy
 is actively used at an attachpoint Status INACTIVE -- Policy is applied by an unattached
policy

      Usage/Status         count
------------------------------------------------------------
      Direct               1
```

```
      Indirect             0

      ACTIVE               0
      INACTIVE             0
      UNUSED               1


      Usage      Status     Route-policy
-------------------------------------------------------------

      Direct     UNUSED     set-rip-unreachable
```

show rpl prefix-set

To display the configuration of prefix sets, use the **show rpl prefix-set** command in EXEC mode.

show rpl prefix-set [*name* | **states** | **brief**]

Syntax Description

name	(Optional) Name of the prefix set.
states	(Optional) Shows all unused, inactive, and active states.
brief	(Optional) Limits the display to a list of the names of all extended community sets without their configurations.

Command Default

No default behavior or values

Command Modes

EXEC

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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Because sets cannot hierarchically reference other sets or policies, no **detail** keyword exists as with the **show rpl policy** command.

Task ID

Task ID	Operations
route-policy	read

Examples

In the following example, the configuration of prefix set pset1 is displayed:

```
RP/0/RSP0/CPU0:router# show rpl prefix-set pset1
!
prefix-set pset1
 10.0.0.1/0,
 10.0.0.2/0 ge 25 le 32,
 10.0.0.5/8 ge 8 le 32,
 10.168.0.0/16 ge 16 le 32,
 172.16.0.9/20 ge 20 le 32,
 192.168.0.5/20 ge 20 le 32
end-set
```

Related Commands

Command	Description
show rpl as-path-set, on page 1805	Displays the configuration of a named AS path set.

Command	Description
show rpl community-set, on page 1812	Displays the configuration of a named community set.
show rpl extcommunity-set, on page 1819	Displays the configuration of a named extended community set.
show rpl route-policy, on page 1856	Displays the configuration of a named route policy.

show rpl prefix-set attachpoints

To display all the policies used at an attach point that reference the named prefix set, use the **show rpl prefix-set attachpoints** command in EXEC mode.

show rpl prefix-set *name* **attachpoints**

Syntax Description

name	Name of a prefix set.

Command Default

No default behavior or values

Command Modes

EXEC

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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **show rpl prefix-set attachpoints** command to display all the policies used at an attach point that reference the named prefix set either directly or indirectly.

The prefix set name is required.

Task ID

Task ID	Operations
route-policy	read

Examples

This example shows the following sample configuration:

```
router bgp 2
 address-family ipv4 unicast
 !
 neighbor 10.0.101.2
  remote-as 100
  address-family ipv4 unicast
   route-policy policy_1 in
  !
 !
 neighbor 10.0.101.3
  remote-as 12
  address-family ipv4 unicast
   route-policy policy_2 in
  !
 !
!
route-policy policy_1
  if (destination in prefix_set_ex1) then
    set local-preference 100
  endif
```

```
  if (as-path in as_path_set_ex1) then
    set community  (10:333) additive
  endif
end-policy
!
route-policy policy_2
  if (destination in prefix_set_ex1) then
    if (community matches-any comm_set_ex1) then
      set community  (10:666) additive
    endif
    if (extcommunity matches-any ext_comm_set_rt_ex1) then
      set community  (10:999) additive
    endif
  endif
end-policy
```

Given this sample configuration, the **show rpl prefix-set prefix_set_ex1 attachpoints** command displays the following information:

```
RP/0/RSP0/CPU0:router# show rpl prefix-set prefix_set_ex1 attachpoints

BGP Attachpoint:Neighbor

Neighbor/Group  type  afi/safi   in/out     referring policy attached policy
-------------------------------------------------------------------------------
10.0.101.2      --    IPv4/uni   in         policy_1         policy_1
10.0.101.3      --    IPv4/uni   in         policy_2         policy_2
```

This table describes the significant fields shown in the display.

Table 180: show rpl prefix-set attachpoints Field Descriptions

Field	Description
BGP Attachpoint	Location of the attach point.
Neighbor/Group	IP address of the attach point on the neighbor.
type	Address family mode.
afi/safi	Address family identifier or subsequent address family identifier.
in/out	Import or export policy.
referring policy	Policy that refers to the AS path set.
attached policy	Policy used at the attach point.

Related Commands

Command	Description
show rpl as-path-set attachpoints, on page 1807	Displays all the policies used at an attach point that reference the named AS path set.
show rpl community-set attachpoints, on page 1814	Displays all the policies used at an attach point that reference the named community set.

Command	Description
show rpl route-policy attachpoints, on page 1859	Displays all the policies used at an attach point that reference the named policy.

show rpl prefix-set references

To list all the policies that reference the named prefix set, use the **show rpl prefix-set references** command in EXEC mode.

show rpl prefix-set *name* **references** [**brief**]

Syntax Description

name	Name of the prefix set.
brief	(Optional) Limits the output to just a summary table and not the detailed information for the named prefix set.

Command Default

No default behavior or values

Command Modes

EXEC

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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **show rpl prefix-set references** command to list all the policies that reference the named prefix set.

Use the optional **brief** keyword to limit the output to just a summary table and not the detailed information for the named prefix set.

Task ID

Task ID	Operations
route-policy	read

Examples

This example shows the following sample configuration:

```
prefix-set ten-net
 10.0.0.0/16 le 32
end-set
prefix-set too-specific
 0.0.0.0/0 ge 25 le 32
end-set
route-policy example-one
 if destination in ten-net then
  drop
 else
  set local-preference 200
  apply set-comms
 endif
end-policy
route-policy set-comms
 set community (10:1234) additive
```

```
end-policy

route-policy example-three
 if destination in too-specific then
  drop
 else
  apply example-one
  pass
 endif
end-policy
```

The following example displays information showing the usage and status of each policy that references the prefix set ten-net. The **brief** keyword limits the display to just a summary table and not the detailed information for the prefix set.

```
RP/0/RSP0/CPU0:router# show rpl prefix-set ten-net references brief

Usage Direct -- Reference occurs in this policy
Usage Indirect -- Reference occurs via an apply statement

Status UNUSED -- Policy is not in use at an attachpoint (unattached)
Status ACTIVE -- Policy is actively used at an attachpoint
Status INACTIVE -- Policy is applied by an unattached policy

     Usage/Status         count
-------------------------------------------------------------
     Direct                 1
     Indirect               1

     ACTIVE                 0
     INACTIVE               1
     UNUSED                 1
```

This table describes the significant fields shown in the display.

Table 181: show rpl prefix-set name references Field Descriptions

Field	Description
Usage/Status	Displays the usage and status of all policies that reference the prefix set.
count	Number of policies that match each usage and status option.

Related Commands

Command	Description
show rpl as-path-set references, on page 1810	Lists all the policies that reference the named AS path set.
show rpl community-set references, on page 1816	Lists all the policies that reference the named community set.
show rpl route-policy references, on page 1864	Lists all the policies that reference the named policy.

show rpl rd-set

To display the configuration of route distinguisher (RD) sets, use the **show rpl rd-set** command in EXEC mode.

show rpl rd-set [*name* | **states** | **brief**]

Syntax Description

name	(Optional) Name of the RD set.
states	(Optional) Shows all unused, inactive, and active states.
brief	(Optional) Limits the display to a list of the names of all RD sets without their configurations.

Command Default

No default behavior or values

Command Modes

EXEC

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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Because sets cannot hierarchically reference other sets or policies, no **detail** keyword exists as with the **show rpl policy** command.

Task ID

Task ID	Operations
route-policy	read

Examples

In the following example, the configuration of RD set rdset1 is displayed:

```
RP/0/RSP0/CPU0:router# show rpl rd-set rdset1

rd-set rdset1
  10:151,
  100.100.100.1:153,
  100.100.100.62/31:63
end-set
```

Related Commands

Command	Description
show rpl as-path-set, on page 1805	Displays the configuration of a named AS path set.
show rpl community-set, on page 1812	Displays the configuration of a named community set.

Command	Description
show rpl extcommunity-set, on page 1819	Displays the configuration of a named extended community set.
show rpl prefix-set, on page 1843	Displays the configuration of a named prefix set.
show rpl route-policy, on page 1856	Displays the configuration of a named route policy.

show rpl rd-set attachpoints

To display all the policies used at an attach point that reference the named route distinguisher (RD) set, use the **show rpl rd-set attachpoints** command in EXEC mode.

show rpl rd-set *name* **attachpoints**

Syntax Description

name	Name of an RD set.

Command Default

No default behavior or values

Command Modes

EXEC

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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **show rpl rd-set attachpoints** command to display all the policies used at an attach point that reference the named RD set either directly or indirectly.

Task ID

Task ID	Operations
route-policy	read

Examples

This example shows the following sample configuration:

```
route-policy rdsetmatch
  if rd in rdset1 then
    set community (10:112)
  elseif rd in rdset2 then
    set community (10:223)
  endif
end-policy

router bgp 10
address-family vpnv4 unicast
 exit
 neighbor 10.0.101.1
  remote-as 11
  address-family vpnv4 unicast
   route-policy rdsetmatch in
!
```

Given this sample configuration, the **show rpl rd-set rdset1 attachpoints** command displays the following information:

```
RP/0/RSP0/CPU0:router# show rpl rd-set rdset attachpoints
```

```
BGP Attachpoint: Neighbor

Neighbor/Group  type  afi/safi   in/out   vrf name
-------------------------------------------------
10.0.101.1      --    IPv4/vpn   in         default
```

This table describes the significant fields shown in the display.

Table 182: show rpl rd-set attachpoints Field Descriptions

Field	Description
Neighbor/Group	BGP neighbor or neighbor group where the specified RD is used.
afi/safi	BGP address family or subaddress family where the RD set is used.
in/out	Direction
vrf name	VRF name where the RD set is used.

Related Commands

Command	Description
show rpl as-path-set attachpoints, on page 1807	Displays all the policies used at an attach point that reference the named AS path set.
show rpl community-set attachpoints, on page 1814	Displays all the policies used at an attach point that reference the named community set.
show rpl prefix-set attachpoints, on page 1845	Displays all the policies used at an attach point that reference the named prefix set.
show rpl route-policy attachpoints, on page 1859	Displays all the policies used at an attach point that reference the named policy.

show rpl rd-set references

To list all the policies that reference the named route distinguisher (RD) set, use the **show rpl rd-set references** command in EXEC mode.

show rpl rd-set *name* **references** [**brief**]

Syntax Description

name	Name of the RD set.
brief	(Optional) Limits the output to just a summary table and not the detailed information for the RD set.

Command Default

No default behavior or values

Command Modes

EXEC

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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **show rpl rd-set references** command to list all the policies that reference the named RD set.

Use the optional **brief** keyword to limit the output to just a summary table and not the detailed information for the named RD set.

Task ID

Task ID	Operations
route-policy	read

Examples

This example shows the following sample configuration:

```
route-policy rdsetmatch
  if rd in rdset1 then
    set community (10:112)
  elseif rd in rdset2 then
    set community (10:223)
  endif
end-policy
!
router bgp 10
 address-family vpnv4 unicast
 !
 neighbor 10.0.101.1
  remote-as 11
  address-family vpnv4 unicast
   route-policy rdsetmatch in
  !
```

Given this sample configuration, the **show rpl rd-set rdset1 references** command displays the following information:

```
RP/0/RSP0/CPU0:router# show rpl rd-set rdset1 references

Usage Direct -- Reference occurs in this policy
Usage Indirect -- Reference occurs via an apply statement

Status UNUSED -- Policy is not in use at an attachpoint (unattached)
Status ACTIVE -- Policy is actively used at an attachpoint
Status INACTIVE -- Policy is applied by an unattached policy

     Usage/Status         count
-----------------------------------------------------------------
     Direct               1
     Indirect             0

     ACTIVE               1
     INACTIVE             0
     UNUSED               0


     route-policy         usage      policy status
-----------------------------------------------------------------
    rdsetmatch           Direct     ACTIVE
```

This table describes the significant fields shown in the display.

Table 183: show rpl rd-set name references Field Descriptions

Field	Description
route-policy	Name of the route policy.
usage	Type of reference usage for the route policy.
policy status	Status of the route policy.

Related Commands

Command	Description
show rpl as-path-set references, on page 1810	Lists all the policies that reference the named AS path set.
show rpl community-set references, on page 1816	Lists all the policies that reference the named community set.
show rpl prefix-set references, on page 1848	Lists all the policies that reference the named prefix set.
show rpl route-policy references, on page 1864	Lists all policies that reference the named policy.

show rpl route-policy

To display the configuration of route policies, use the **show rpl route-policy** command in EXEC mode.

show rpl route-policy [*name* [**detail**] | **states** | **brief**]

Syntax Description

name	(Optional) Name of a route policy.
detail	(Optional) Displays the configuration of all policies and sets that a policy uses.
states	(Optional) Shows all unused, inactive, and active states.
brief	(Optional) Limits the display to a list of the names of all extended community sets without their configurations.

Command Default

No default behavior or values

Command Modes

EXEC

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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the optional **brief** keyword to limit the display to a list of the names of policies without their configurations.

Task ID

Task ID	Operations
route-policy	read

Examples

In the following example, the configuration of a route policy named policy_1 is displayed.

```
RP/0/RSP0/CPU0:router# show rpl route-policy policy_1

route-policy policy_1
  if destination in prefix_set_1 and not destination in sample1 then
    if as-path in aspath_set_1 then
      set local-preference 300
      set origin igp
    elseif as-path in as_allowed then
      set local-preference 400
      set origin igp
    else
      set origin igp
    endif
  else
    drop
```

```
  endif
  set med 120
  set community (8660:612) additive
  apply set_lpref_from_comm
end-policy
```

If the optional **detail** keyword is used, all routing policy language (RPL) policies and sets that route policy policy_1 uses are displayed, as shown in the following example.

```
RP/0/RSP0/CPU0:router# show rpl route-policy policy_1 detail

!
prefix-set sample1
  0.0.0.0/0,
  0.0.0.0/0 ge 25 le 32,
  10.0.0.0/8 ge 8 le 32,
  192.168.0.0/16 ge 16 le 32,
  224.0.0.0/20 ge 20 le 32,
  240.0.0.0/20 ge 20 le 32
end-set
!
prefix-set prefix_set_1
 10.0.0.1/24 ge 24 le 32,
 10.0.0.5/24 ge 24 le 32,
 172.16.0.1/24 ge 24 le 32,
 172.16.5.5/24 ge 24 le 32,
 172.16.20.10/24 ge 24 le 32,
 172.30.0.1/24 ge 24 le 32,
 10.0.20.10/24 ge 24 le 32,
 172.18.0.5/24 ge 24 le 32,
 192.168.0.1/24 ge 24 le 32,
 192.168.20.10/24 ge 24 le 32,
 192.168.200.10/24 ge 24 le 32,
 192.168.255.254/24 ge 24 le 32
end-set
!
as-path-set as_allowed
  ios-regex '.* _1239_ .*',
  ios-regex '.* _3561_ .*',
  ios-regex '.* _701_ .*',
  ios-regex '.* _666_ .*',
  ios-regex '.* _1755_ .*',
  ios-regex '.* _1756_ .*'
end-set
!
as-path-set aspath_set_1
  ios-regex '_9148_',
  ios-regex '_5870_',
  ios-regex '_2408_',
  ios-regex '_2531_',
  ios-regex '_197_',
  ios-regex '_2992_'
end-set
!
route-policy set_lpref_from_comm
  if community matches-any (2:50) then
    set local-preference 50
  elseif community matches-any (2:60) then
    set local-preference 60
  elseif community matches-any (2:70) then
    set local-preference 70
  elseif community matches-any (2:80) then
```

```
    set local-preference 80
  elseif community matches-any (2:90) then
    set local-preference 90
  endif
end-policy
!
route-policy policy_1
  if destination in prefix_set_1 and not destination in sample1 then
    if as-path in aspath_set_1 then
      set local-preference 300
      set origin igp
    elseif as-path in as_allowed then
      set local-preference 400
      set origin igp
    else
      set origin igp
    endif
  else
    drop
  endif
  set med 120
  set community (8660:612) additive
  apply set_lpref_from_comm
end-policy
```

Related Commands

Command	Description
show rpl as-path-set, on page 1805	Displays the configuration of a named AS path set.
show rpl community-set, on page 1812	Displays the configuration of a named community set.
show rpl extcommunity-set, on page 1819	Displays the configuration of a named extended community set.
show rpl prefix-set, on page 1843	Displays the configuration of a named prefix set.

show rpl route-policy attachpoints

To display all the policies used at an attach point that reference the named policy, use the **show rpl route-policy attachpoints** command in EXEC mode.

show rpl route-policy *name* **attachpoints**

Syntax Description

name	Name of a policy.

Command Default

No default behavior or values

Command Modes

EXEC

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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **show rpl route-policy attachpoints** command to display all the policies used at an attach point that reference the named policy either directly or indirectly.

The policy name is required.

Task ID

Task ID	Operations
route-policy	read

Examples

This example shows the following sample configuration:

```
router bgp 2
 address-family ipv4 unicast
 !
 neighbor 10.0.101.2
  remote-as 100
  address-family ipv4 unicast
   route-policy policy_1 in
  !
 !
 neighbor 10.0.101.3
  remote-as 12
  address-family ipv4 unicast
   route-policy policy_2 in
  !
 !
!

RP/0/RSP0/CPU0:router# show rpl route-policy policy_1
```

```
route-policy policy_1
  if (destination in prefix_set_ex1) then
    set local-preference 100
  endif
  if (as-path in as_path_set_ex1) then
    set community  (10:333) additive
  endif
end-policy
!
RP/0/RSP0/CPU0:router# show rpl route-policy policy_2

route-policy policy_2
  if (destination in prefix_set_ex1) then
    if (community matches-any comm_set_ex1) then
      set community  (10:666) additive
    endif
    if (extcommunity matches-any ext_comm_set_rt_ex1) then
      set community  (10:999) additive
    endif
  endif
end-policy
!
```

The following command displays the route policy attach points for policy_2:

```
RP/0/RSP0/CPU0:router# show rpl route-policy policy_2 attachpoints

BGP Attachpoint: Neighbor

Neighbor/Group  type  afi/safi   in/out   vrf name
-------------------------------------------------
10.0.101.2      --    IPv4/uni   in         default
10.0.101.2      --    IPv4/uni   out        default
```

This table describes the significant fields shown in the display.

Table 184: show rpl route-policy attachpoints Field Descriptions

Field	Description
BGP Attachpoint	Location of the attach point.
Neighbor/Group	IP address of the attach point on the neighbor.
type	Displays the address family mode.
afi/safi	Address family identifier or subsequent address family identifier.
vrf name	Name of the VPN routing and forwarding (VRF) instance.

Related Commands

Command	Description
show rpl as-path-set attachpoints, on page 1807	Displays all the policies used at an attach point that reference the named AS path set.

Command	Description
show rpl community-set attachpoints, on page 1814	Displays all the policies used at an attach point that reference the named community set.
show rpl prefix-set attachpoints, on page 1845	Displays all the policies used at an attach point that reference the named prefix set.

show rpl route-policy inline

To display all policies and sets that a policy uses expanded inline, use the **show rpl route-policy inline** command in EXEC mode.

show rpl route-policy *name* **inline**

Syntax Description

name	Name of a policy.

Command Default

No default behavior or values

Command Modes

EXEC

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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **show rpl route-policy inline** command to examine the configuration of a specified route policy. All policies and sets that a policy uses are gathered together and displayed expanded inline.

The policy name is required.

Task ID

Task ID	Operations
route-policy	read

Examples

The following command displays the route policy policy_1:

```
RP/0/RSP0/CPU0:router# show rpl route-policy policy_1

!
route-policy policy_1
  if destination in prefix_set_1 and not destination in martians then
    if as-path in aspath_set_1 then
      set local-preference 300
      set origin igp
    elseif as-path in as_allowed then
      set local-preference 400
      set origin igp
    else
      set origin igp
    endif
  else
    drop
  endif
  set med 120
  set community (8660:612) additive
```

```
  apply set_lpref_from_comm
end-policy
```

The following command displays the route policy policy_1 and all the other sets or policies it refers too inline. Adding the inline keyword causes the configuration to be displayed inline for all RPL objects that the route-policy policy_1 uses.

```
RP/0/RSP0/CPU0:router#show rpl policy policy_1 inline

route-policy policy_1
  if destination in (91.5.152.0/24 ge 24 le 32, 91.220.152.0/24 ge 24 le 32, 61.106.52.0/24
 ge 24 le 32, 222.168.199.0/24
  ge 24 le 32, 93.76.114.0/24 ge 24 le 32, 41.195.116.0/24 ge 24 le 32, 35.92.152.0/24 ge
24 le 32, 143.144.96.0/24 ge 24
  le 32, 79.218.81.0/24 ge 24 le 32, 75.213.219.0/24 ge 24 le 32, 178.220.61.0/24 ge 24 le
 32, 27.195.65.0/24 ge 24 le 32)
  and  not destination in (0.0.0.0/0, 0.0.0.0/0 ge 25 le 32, 10.0.0.0/8 ge 8 le 32,
192.168.0.0/16 ge 16 le 32, 224.0.0.0/20
  ge 20 le 32, 240.0.0.0/20 ge 20 le 32) then
    if as-path in (ios-regex '_9148_', ios-regex '_5870_', ios-regex '_2408_', ios-regex
'_2531_', ios-regex '_197_',
 ios-regex '_2992_') then
      set local-preference 300
      set origin igp
    elseif as-path in
 (ios-regex '.* _1239_ .*', ios-regex '.* _3561_ .*', ios-regex '.* _701_ .*', ios-regex
'.* _666_ .*', ios-regex '.* _1755_ .*',
 ios-regex '.* _1756_ .*') then
      set local-preference 400
      set origin igp
    else
      set origin igp
    endif
  else
    drop
  endif
  set med 120
  set community (8660:612) additive
  # apply set_lpref_from_comm
  if community matches-any (2:50) then
    set local-preference 50
  elseif community matches-any (2:60) then
    set local-preference 60
  elseif community matches-any (2:70) then
    set local-preference 70
  elseif community matches-any (2:80) then
    set local-preference 80
  elseif community matches-any (2:90) then
    set local-preference 90
  endif
  # end-apply set_lpref_from_comm
end-policy
```

show rpl route-policy references

To list all the policies that reference the named policy, use the **show rpl route-policy references** command in EXEC mode.

show rpl route-policy *name* **references** [**brief**]

Syntax Description

name	Name of a prefix set.
brief	(Optional) Limits the output to just a summary table and not the detailed information for the named policy.

Command Default

No default behavior or values

Command Modes

EXEC

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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **show rpl route-policy references** command to list all the policies that reference the named policy.

Use the optional **brief** keyword to limit the output to just a summary table and not the detailed information for the policy.

Task ID

Task ID	Operations
route-policy	read

Examples

This example shows the following sample configuration:

```
prefix-set ten-net
 10.0.0.0/16 le 32
end-set
prefix-set too-specific
 0.0.0.0/0 ge 25 le 32
end-set
route-policy example-one
 if destination in ten-net then
  drop
 else
  set local-preference 200
  apply set-comms
 endif
end-policy
route-policy set-comms
 set community (10:1234) additive
```

```
end-policy
route-policy example-three
 if destination in too-specific then
  drop
 else
  apply example-one
  pass
 endif
end-policy
```

The following command displays information about the policy set-comms and how it is referenced:

```
RP/0/RSP0/CPU0:router# show rpl route-policy set-comms references

Usage Direct -- Reference occurs in this policy
Usage Indirect -- Reference occurs via an apply statement

Status UNUSED -- Policy is not in use at an attachpoint (unattached)
Status ACTIVE -- Policy is actively used at an attachpoint
Status INACTIVE -- Policy is applied by an unattached policy

     Usage/Status         count
-------------------------------------------------------------
     Direct               1
     Indirect             1

     ACTIVE               0
     INACTIVE             1
     UNUSED               1

     route-policy         usage  policy status
-------------------------------------------------------------
    example-one          Direct  INACTIVE
    example-three        Indirect UNUSED
```

The direct usage indicates that the route policy example-one directly applies the policy set-comms, that is, example-one has a line in the form apply set-comms. The usage Indirect indicates that the route policy example-three does not directly apply the route policy set-comms. However, the route policy example-three does apply the policy example-one, which in turn applies the policy set-comms, so there is an indirect reference from example-three to the route policy set-comms.

The status column indicates one of three states. A policy is active if it is in use at an attach point. In the example provided, neither example-one nor example-three is in use at an attach point, which leaves two possible states: UNUSED or INACTIVE. The route policy example-one is inactive because it has some other policy (example-three) that references it, but neither example-one nor any of the policies that reference it (example-one) are in use at an attach point. The route policy example-three has a status of unused because it is not used at an attach point and no other route policies in the system refer to it.

This table describes the significant fields shown in the display.

Table 185: show rpl route-policy references Field Descriptions

Field	Description
Usage/Status	Displays the usage and status of all policies that reference the specified policy. Values for usage are Direct or Indirect. Values for status are ACTIVE, INACTIVE, and UNUSED.
count	Number of policies that match each usage and status option.
route-policy	One name for multiple policies that reference the specified policy.
usage	Type of usage for the policy.
policy status	Status of the policy.

Related Commands

Command	Description
show rpl as-path-set references, on page 1810	Lists all policies that reference the named AS path set.
show rpl community-set references, on page 1816	Lists all policies that reference the named community set.
show rpl prefix-set references, on page 1848	Lists all policies that reference the named prefix set.

show rpl route-policy uses

To display information about a specified named policy, use the **show rpl route-policy uses** command in EXEC mode.

show rpl route-policy *name* **uses** {**policies** | **sets** | **all**} [**direct**]

Syntax Description

name	Name of a policy.
policies	Generates a list of all policies that the named policy uses.
sets	Lists all named sets that are used by the policy.
all	Generates a list of both sets and policies that the named policy references.
direct	(Optional) Lists only the policies or sets used directly in the named policy block. Set or policy references that occur as a result of an **apply** statement are not listed.

Command Default

No default behavior or values

Command Modes

EXEC

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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **show rpl route-policy uses** command to display information about a specified named policy.

Task ID

Task ID	Operations
route-policy	read

Examples

This example shows the following sample configuration:

```
prefix-set ten-net
 10.0.0.0/16 le 32
end-set
prefix-set too-specific
 0.0.0.0/0 ge 25 le 32
end-set
route-policy example-one
 if destination in ten-net then
  drop
 else
  set local-preference 200
  apply set-comms
 endif
```

```
end-policy
route-policy set-comms
 set community (10:1234) additive
end-policy
route-policy example-three
 if destination in too-specific then
  drop
 else
  apply example-one
  pass
 endif
end-policy
```

The following command lists the policies one and set-comms. It also lists the prefix sets too-specific and ten-net.

```
RP/0/RSP0/CPU0:router# show rpl route-policy example-three uses all

Policies directly and indirectly applied by this policy:
---------------------------------------------------------
    example-one set-comms

Sets referenced directly and indirectly
 ----------------------------------------
(via applied policies) in this policy:

type prefix-set:
    ten-net too-specific
```

The sets example-one and set-comms are listed as policies that are used by the policy example-three. The policy example-one is listed because route policy example-three uses it in an **apply** statement. The policy set-comms is also listed because example-one applies it. Similarly, the prefix-set too-specific is used directly in the **if** statement in the policy example-three, and the prefix-set ten-net is used in the policy example-one. The optional **direct** keyword can be used to limit the output to just those sets and policies that are used within the example-three block itself, as shown in the following example:

```
RP/0/RSP0/CPU0:router# show rpl route-policy example-three uses all direct

Policies directly applied by this policy:
-----------------------------------------
    example-one

Sets used directly in this policy
-----------------------------------
type prefix-set:
    too-specific
```

As can be seen in the output, the route policy set-comms and the prefix set ten-net are no longer included in the output when the **direct** keyword is used. The **direct** form of the command considers only those sets or policies used in the specified route policy and any additional policies or sets that may be used if you follow the hierarchy of **apply** statements.

This table describes the significant fields shown in the display.

Table 186: show rpl route-policy uses Field Descriptions

Field	Description
type	Displays the type used in the policy configuration. Values for type are prefix-set, community-set, extcommunity-set, and as-path-set.

show rpl unused as-path-set

To display the AS path sets that are defined but not used by a policy at an attach point or referenced in a policy using an **apply** statement, use the **show rpl unused as-path-set** command in EXEC mode.

show rpl unused as-path-set [**detail**]

Syntax Description

detail	(Optional) Displays the content of the object and all referenced objects for unused AS path sets.

Command Default

No default behavior or values

Command Modes

EXEC

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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **show rpl unused as-path-set** command to display all AS path sets that are not used in a policy at an attach point either directly or indirectly and are not referenced by any policies in the system.

Task ID

Task ID	Operations
route-policy	read

Examples

This example shows the following sample configuration:

```
router bgp 2
 address-family ipv4 unicast
 !
 neighbor 10.0.101.2
  remote-as 100
  address-family ipv4 unicast
   route-policy policy_1 in
  !
 !
 neighbor 10.0.101.3
  remote-as 12
  address-family ipv4 unicast
   route-policy policy_2 in
  !
 !
!
as-path-set as_path_set_ex1
  ios-regex '^_65500_$',
  ios-regex '^_65501_$'
end-set
!
as-path-set as_path_set_ex2
```

```
  ios-regex '^_65502_$',
  ios-regex '^_65503_$'
end-set
!
as-path-set as_path_set_ex3
  ios-regex '^_65504_$',
  ios-regex '^_65505_$'
end-set
!
route-policy sample
  if (destination in sample) then
    drop
  endif
end-policy
!
route-policy policy_1
  if (destination in prefix_set_ex1) then
    set local-preference 100
  endif
  if (as-path in as_path_set_ex1) then
    set community  (10:333) additive
  endif
end-policy
!
route-policy policy_2
  if (destination in prefix_set_ex1) then
    if (community matches-any comm_set_ex1) then
      set community  (10:666) additive
    endif
    if (extcommunity matches-any ext_comm_set_rt_ex1) then
      set community  (10:999) additive
    endif
  endif
end-policy
!
route-policy policy_3
  if (destination in prefix_set_ex2) then
    set local-preference 100
  endif
  if (as-path in as_path_set_ex2) then
    set community  (10:333) additive
  endif
end-policy
!
route-policy policy_4
  if (destination in prefix_set_ex2) then
    if (community matches-any comm_set_ex2) then
      set community  (10:666) additive
    endif
    if (extcommunity matches-any ext_comm_set_rt_ex2) then
      set community  (10:999) additive
    endif
  endif
end-policy
!
route-policy policy_5
  apply sample
  apply policy_3
end-policy
```

Given this sample configuration, the **show rpl unused as-path-set** command displays the following information:

```
RP/0/RSP0/CPU0:router# show rpl unused as-path-set

ACTIVE -- Referenced by at least one policy which is attached
INACTIVE -- Only referenced by policies which are not attached
UNUSED -- Not attached (directly or indirectly) and not referenced

The following as-path-sets are UNUSED
-------------------------------------
as_path_set_ex3
```

Related Commands

show rpl unused community-set, on page 1873	Displays the community sets that are not referenced at all.
show rpl unused extcommunity-set, on page 1876	Displays the extended community sets that are not referenced at all.
show rpl unused prefix-set, on page 1878	Displays the prefix sets that are not referenced at all.
show rpl unused rd-set, on page 1881	Displays the RD sets that are not referenced at all.
show rpl unused route-policy, on page 1883	Displays the route policies that are not referenced at all.

show rpl unused community-set

To display the community sets that are defined but not used by a policy at an attach point or referenced in a policy using an **apply** statement, use the **show rpl unused community-set** command in EXEC mode.

show rpl unused community-set [**detail**]

Syntax Description

detail	(Optional) Displays the content of the object and all referenced objects for unused community sets.

Command Default

No default behavior or values

Command Modes

EXEC

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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **show rpl unused community-set** command to display all the community sets that are not used in a policy at an attach point either directly or indirectly and are not referenced by any policies in the system.

Task ID

Task ID	Operations
route-policy	read

Examples

This example shows the following sample configuration:

```
router bgp 2
 address-family ipv4 unicast
 !
 neighbor 10.0.101.2
  remote-as 100
  address-family ipv4 unicast
   route-policy policy_1 in
  !
 !
 neighbor 10.0.101.3
  remote-as 12
  address-family ipv4 unicast
   route-policy policy_2 in
  !
 !
!

community-set comm_set_ex1
  65500:1,
  65500:2,
  65500:3
end-set
```

```
!
community-set comm_set_ex2
  65501:1,
  65501:2,
  65501:3
end-set
!
community-set comm_set_ex3
  65502:1,
  65502:2,
  65502:3
end-set
!
route-policy sample
  if (destination in sample) then
    drop
  endif
end-policy
!
route-policy policy_1
  if (destination in prefix_set_ex1) then
    set local-preference 100
  endif
  if (as-path in as_path_set_ex1) then
    set community  (10:333) additive
  endif
end-policy
!
route-policy policy_2
  if (destination in prefix_set_ex1) then
    if (community matches-any comm_set_ex1) then
      set community  (10:666) additive
    endif
    if (extcommunity matches-any ext_comm_set_rt_ex1) then
      set community  (10:999) additive
    endif
  endif
end-policy
!
route-policy policy_3
  if (destination in prefix_set_ex2) then
    set local-preference 100
  endif
  if (as-path in as_path_set_ex2) then
    set community  (10:333) additive
  endif
end-policy
!
route-policy policy_4
  if (destination in prefix_set_ex2) then
    if (community matches-any comm_set_ex2) then
      set community  (10:666) additive
    endif
    if (extcommunity matches-any ext_comm_set_rt_ex2) then
      set community  (10:999) additive
    endif
  endif
end-policy
!
route-policy policy_5
  apply sample
  apply policy_3
end-policy
```

Given this sample configuration, the **show rpl unused community-set** command displays the following information:

```
RP/0/RSP0/CPU0:router# show rpl unused community-set

ACTIVE -- Referenced by at least one policy which is attached
INACTIVE -- Only referenced by policies which are not attached
UNUSED -- Not attached (directly or indirectly) and not referenced

The following community-sets are UNUSED
----------------------------------------
comm_set_ex3
```

Related Commands

Command	Description
show rpl unused as-path-set, on page 1870	Displays the AS path sets that are not referenced at all.
show rpl unused extcommunity-set, on page 1876	Displays the extended community sets that are not referenced at all.
show rpl unused prefix-set, on page 1878	Displays the prefix sets that are not referenced at all.
show rpl unused rd-set, on page 1881	Displays the RD sets that are not referenced at all.
show rpl unused route-policy, on page 1883	Displays the route policies that are not referenced at all.

show rpl unused extcommunity-set

To display the extended community sets that are defined but not used by a policy at an attach point or referenced in a policy using an **apply** statement, use the **show rpl unused extcommunity-set** command in EXEC mode.

show rpl unused extcommunity-set [**cost** | **detail** | **rt** | **soo**]

Syntax Description

cost	(Optional) Displays the unused extended-community cost objects.
rt	(Optional) Displays the unused extended community RT objects.
soo	(Optional) Displays the unused extended-community SoO objects.
detail	(Optional) Displays the content of the object and all referenced objects for unused extended community sets.

Command Default

No default behavior or values

Command Modes

EXEC

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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **show rpl unused extcommunity-set** command to display all extended community sets that are not used in a policy at an attach point either directly or indirectly and are not referenced by any policies in the system.

Task ID

Task ID	Operations
route-policy	read

Examples

The following is sample output for the **show rpl unused extcommunity-set** command:

```
RP/0/RSP0/CPU0:router:router# show rpl unused extcommunity-set

ACTIVE -- Referenced by at least one policy which is attached
INACTIVE -- Only referenced by policies which are not attached
UNUSED -- Not attached (directly or indirectly) and not referenced

The following extcommunity-sets are UNUSED
-------------------------------------------
ext_comm_set_ex3
```

Related Commands

Command	Description
show rpl unused as-path-set, on page 1870	Displays the AS path sets that are not referenced at all.
show rpl unused community-set, on page 1873	Displays community sets that are not referenced at all.
show rpl unused prefix-set, on page 1878	Displays prefix sets that are not referenced at all.
show rpl unused rd-set, on page 1881	Displays the RD sets that are not referenced at all.
show rpl unused route-policy, on page 1883	Displays the route policies that are not referenced at all.

show rpl unused prefix-set

To display the prefix sets that are defined but not used by a policy at an attach point or referenced in a policy using an **apply** statement, use the **show rpl unused prefix-set** command in EXEC mode.

show rpl unused prefix-set [**detail**]

Syntax Description

detail	(Optional) Displays the content of the object and all referenced objects for unused prefix sets.

Command Default

No default behavior or values

Command Modes

EXEC

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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **show rpl unused prefix-set** command to display all prefix sets that are not used in a policy at an attach point either directly or indirectly and are not referenced by any policies in the system.

Task ID

Task ID	Operations
route-policy	read

Examples

This example shows the following sample configuration:

```
router bgp 2
 address-family ipv4 unicast
 !
 neighbor 10.0.101.2
  remote-as 100
  address-family ipv4 unicast
   route-policy policy_1 in
  !
 !
 neighbor 10.0.101.3
  remote-as 12
  address-family ipv4 unicast
   route-policy policy_2 in
  !
 !
!

prefix-set sample
  0.0.0.0/0,
  0.0.0.0/0 ge 25 le 32,
  10.0.0.0/8 ge 8 le 32,
  192.168.0.0/16 ge 16 le 32,
```

```
  224.0.0.0/20 ge 20 le 32,
  240.0.0.0/20 ge 20 le 32
end-set
!
prefix-set prefix_set_ex1
  10.0.0.0/16 ge 16 le 32,
  0.0.0.0/0 ge 25 le 32,
  0.0.0.0/0
end-set
!
prefix-set prefix_set_ex2
  220.220.220.0/24 ge 24 le 32,
  220.220.120.0/24 ge 24 le 32,
  220.220.130.0/24 ge 24 le 32
end-set
!
prefix-set prefix_set_ex3
  221.221.220.0/24 ge 24 le 32,
  221.221.120.0/24 ge 24 le 32,
  221.221.130.0/24 ge 24 le 32
end-set
!
route-policy sample
  if (destination in sample) then
    drop
  endif
end-policy
!
route-policy policy_1
  if (destination in prefix_set_ex1) then
    set local-preference 100
  endif
  if (as-path in as_path_set_ex1) then
    set community  (10:333) additive
  endif
end-policy
!
route-policy policy_2
  if (destination in prefix_set_ex1) then
    if (community matches-any comm_set_ex1) then
      set community  (10:666) additive
    endif
    if (extcommunity matches-any ext_comm_set_rt_ex1) then
      set community  (10:999) additive
    endif
  endif
end-policy
!
route-policy policy_3
  if (destination in prefix_set_ex2) then
    set local-preference 100
  endif
  if (as-path in as_path_set_ex2) then
    set community  (10:333) additive
  endif
end-policy
!
route-policy policy_4
  if (destination in prefix_set_ex2) then
    if (community matches-any comm_set_ex2) then
      set community  (10:666) additive
    endif
    if (extcommunity matches-any ext_comm_set_rt_ex2) then
      set community  (10:999) additive
```

```
    endif
  endif
end-policy
!
route-policy policy_5
  apply sample
  apply policy_3
end-policy
--------------------------
ext_comm_set_ex3
```

Given this sample configuration, the **show rpl unused prefix-set** command displays the following information:

```
RP/0/RSP0/CPU0:router# show rpl unused prefix-set

ACTIVE -- Referenced by at least one policy which is attached
INACTIVE -- Only referenced by policies which are not attached
UNUSED -- Not attached (directly or indirectly) and not referenced

The following prefix-sets are UNUSED
-------------------------------------
prefix_set_ex3
```

Related Commands

Command	Description
show rpl unused as-path-set, on page 1870	Displays AS path sets that are not referenced at all.
show rpl unused community-set, on page 1873	Displays community sets that are not referenced at all.
show rpl unused extcommunity-set, on page 1876	Displays extended community sets that are not referenced at all.
show rpl unused rd-set, on page 1881	Displays the RD sets that are not referenced at all.
show rpl unused route-policy, on page 1883	Displays the route policies that are not referenced at all.

show rpl unused rd-set

To display the route distinguisher (RD) sets that are defined but not used by a policy at an attach point or referenced in a policy using an **apply** statement, use the **show rpl unused rd-set** command in EXEC mode.

show rpl unused rd-set [**detail**]

Syntax Description

detail	(Optional) Displays the content of the object and all referenced objects for unused RD sets.

Command Default

No default behavior or values

Command Modes

EXEC

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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **show rpl unused rd-set** command to display all of the RD sets that are not used in a policy at an attach point either directly or indirectly and are not referenced by any policies in the system.

Task ID

Task ID	Operations
route-policy	read

Examples

The **show rpl unused rd-set** command displays the following information:

```
RP/0/RSP0/CPU0:router# show rpl unused rd-set

ACTIVE -- Referenced by at least one policy which is attached
INACTIVE -- Only referenced by policies which are not attached
UNUSED -- Not attached (directly or indirectly) and not referenced

The following rd-sets are UNUSED
----------------------------------------
None found with this status.
```

Related Commands

Command	Description
show rpl unused as-path-set, on page 1870	Displays the AS path sets that are not referenced at all.
show rpl unused community-set, on page 1873	Displays the community sets that are not referenced at all.
show rpl unused extcommunity-set, on page 1876	Displays the extended community sets that are not referenced at all.

Command	Description
show rpl unused prefix-set, on page 1878	Displays the prefix sets that are not referenced at all.
show rpl unused route-policy, on page 1883	Displays the route policies that are not referenced at all.

show rpl unused route-policy

To display the route policies that are defined but not used at an attach point or referenced using an **apply** statement, use the **show rpl unused route-policy** command in EXEC mode.

show rpl unused route-policy [**detail**]

Syntax Description

detail	(Optional) Displays the content of the object and all referenced objects for unused route policies.

Command Default

No default behavior or values

Command Modes

EXEC

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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **show rpl unused route-policy** command to display route policies that are defined but not used at an attach point or referenced from another policy using an **apply** statement.

Task ID

Task ID	Operations
route-policy	read

Examples

This example shows the following sample configuration:

```
RP/0/RSP0/CPU0:router# show run | begin prefix-set

Building configuration...
prefix-set prefix_set_ex1
  10.0.0.0/16 ge 16 le 32,
  0.0.0.0/0 ge 25 le 32,
  0.0.0.0/0
end-set
!
prefix-set prefix_set_ex2
  220.220.220.0/24 ge 24 le 32,
  220.220.120.0/24 ge 24 le 32,
  220.220.130.0/24 ge 24 le 32
end-set
!
as-path-set as_path_set_ex1
  ios-regex '^_65500_$',
  ios-regex '^_65501_$'
end-set
!
as-path-set as_path_set_ex2
  ios-regex '^_65502_$',
```

```
  ios-regex '^_65503_$'
end-set
!
as-path-set as_path_set_ex3
  ios-regex '^_65504_$',
  ios-regex '^_65505_$'
end-set
!
community-set comm_set_ex1
  65500:1,
  65500:2,
  65500:3
end-set
!
community-set comm_set_ex2
  65501:1,
  65501:2,
  65501:3
end-set
!
extcommunity-set rt ext_comm_set_rt_ex1
  1.2.3.4:34
end-set
!
extcommunity-set rt ext_comm_set_rt_ex2
  2.3.4.5:36
end-set
!
route-policy sample
  if (destination in sample) then
    drop
  endif
end-policy
!
route-policy policy_1
  if (destination in prefix_set_ex1) then
    set local-preference 100
  endif
  if (as-path in as_path_set_ex1) then
    set community (10:333) additive
  endif
end-policy
!
route-policy policy_2
  if (destination in prefix_set_ex1) then
    if (community matches-any comm_set_ex1) then
      set community (10:666) additive
    endif
    if (extcommunity rt matches-any ext_comm_set_rt_ex1) then
      set community (10:999) additive
    endif
  endif
end-policy
!
route-policy policy_3
  if (destination in prefix_set_ex2) then
    set local-preference 100
  endif
  if (as-path in as_path_set_ex2) then
    set community (10:333) additive
  endif
end-policy
!
route-policy policy_4
```

```
  if (destination in prefix_set_ex2) then
    if (community matches-any comm_set_ex2) then
      set community (10:666) additive
    endif
    if (extcommunity rt matches-any ext_comm_set_rt_ex2) then
      set community (10:999) additive
    endif
  endif
end-policy
!
route-policy policy_5
  apply sample
  apply policy_3
end-policy
!
route ipv4 0.0.0.0/0 10.91.37.129
route ipv4 10.91.36.0/23 10.91.37.129
route ipv4 10.91.38.0/24 10.91.37.129
end
```

In the following example, route policies that are defined but not used at an attach point or referenced from another policy using an **apply** statement are displayed using the **show rpl unused route-policy** command.

```
RP/0/RSP0/CPU0:router# show rpl unused route-policy

ACTIVE -- Referenced by at least one policy which is attached
INACTIVE -- Only referenced by policies which are not attached
UNUSED -- Not attached (directly or indirectly) and not referenced

The following policies are (UNUSED)
-----------------------------------
policy_1
policy_2
policy_4
policy_5
```

Related Commands

Command	Description
show rpl unused as-path-set, on page 1870	Displays AS path sets that are not referenced at all.
show rpl unused community-set, on page 1873	Displays community sets that are not referenced at all.
show rpl unused extcommunity-set, on page 1876	Displays extended community sets that are not referenced at all.
show rpl unused prefix-set, on page 1878	Displays prefix sets that are not referenced at all.
show rpl unused rd-set, on page 1881	Displays the RD sets that are not referenced at all.

source in

To test the source of a Border Gateway Protocol (BGP) route against the address contained in either a named or an inline prefix set, use the **source in** command in route-policy configuration mode.

source in {*prefix-set-nameinline-prefix-setparameter*}

Syntax Description

prefix-set-name	Name of a prefix set.
inline-prefix-set	Inline prefix set. The inline prefix set must be enclosed in parentheses.
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 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **source in** command as a conditional expression within an **if** statement to test the source of the route against the data in either a named or an inline prefix set. A comparison that references a prefix set with zero elements in it returns false.

Note For a list of all conditional expressions available within an **if** statement, see the **if** command.

The source of a BGP route is the IP peering address of the neighboring router from which the route was received.

The prefix set can contain both IPv4 and IPv6 prefix specifications.

Task ID

Task ID	Operations
route-policy	read, write

Examples

In the following example, the source of a BGP route is tested against the data in the prefix set my-prefix-set:

```
RP/0/RSP0/CPU0:router(config)# route-policy policy-A
RP/0/RSP0/CPU0:router(config-rpl)# if source in my-prefix-set then
```

In this example, the source of a BGP route is tested against the data in an inline IPv4 prefix set:

```
RP/0/RSP0/CPU0:router(config)# route-policy policy-B
RP/0/RSP0/CPU0:router(config-rpl)# if source in (10.0.0.8, 10.0.0.20) then
```

In this example, the source of a route is tested against the data in an inline IPv6 prefix set:

```
RP/0/RSP0/CPU0:router(config)# route-policy policy-C
RP/0/RSP0/CPU0:router(config-rpl)# if source in (2001:0:0:1::/64, 2001:0:0:2::/64) then
```

Related Commands

Command	Description
prefix-set, on page 1714	Enters a prefix set configuration mode and defines a prefix set.

source rt

To perform import operation based solely on import route-policy by matching on route-targets (RTs) and other criteria specified within the policy, use the **source rt** command in BGP address-family VRF all mode.

source rt import-policy

Syntax Description

This command has no arguments or keywords.

Command Modes

BGP address-family VRF all

Command History

Release	Modification
Release 4.3.0	This command was introduced.

Usage Guidelines

When you configure the **source rt** command, you need not explicitly configure import RTs under global VRF-address family configuration mode. If the import RTs and import route-policy are already defined, then the routes will be imported from RTs configured under import RT and then follows the route-policy attached at import route-policy. In other words, if the import RT is already defined, it will still add the RTs mentioned in the policy to the imported route-targets list but without the use of the **import route-policy** command

Examples

The example shows how to configure the **source rt** command:

```
Router(config)# router bgp 1
Router(config-bgp)# address-family vpnv4 unicast
Router(config-bgp-af)# vrf all
Router(config-bgp-af-vrf-all)# source rt import-policy
```

suppress-route

To indicate that a given component of a BGP aggregate should be suppressed, use the **suppress-route** command in route-policy configuration mode.

suppress-route

Syntax Description

This command has no arguments or keywords.

Command Default

No default behavior or values

Command Modes

Route-policy 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **suppress-route** command to indicate that a given component of an aggregate should be suppressed, that is, not advertised by BGP. See the unsuppress-route, on page 1894 command for information on overriding the **suppress-route** command for individual neighbors.

The **suppress-route** command can be used as an action statement within an **if** statement. For a list of all action statements available within an **if** statement, see the **if** command.

Task ID

Task ID	Operations
route-policy	read, write

Examples

In the following example, if the destination is in 10.1.0.0/16, then the route is not advertised:

```
RP/0/RSP0/CPU0:router(config)# oute-policy check-aggregater
RP/0/RSP0/CPU0:router(config-rpl)# if destination in (10.1.0.0/16) then
RP/0/RSP0/CPU0:router(config-rpl-if)# suppress-route

RP/0/RSP0/CPU0:router(config-rpl-if)# endif
RP/0/RSP0/CPU0:router(config-rpl-if)# end-policy
```

Related Commands

Command	Description
unsuppress-route, on page 1894	Indicates that a given component of an aggregate should be unsuppressed

tag

To match a specific tag value, use the **tag** command in route-policy configuration mode.

tag {**eq** | **ge** | **le** | **is**} {*integerparameter*}

Syntax Description

eq | **ge** | **le** | **is**	Equal to; greater than or equal to; less than or equal to.
integer	Integer value. Range is 0 to 4294967295.
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 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **tag** command as a conditional expression within an **if** statement to match a specific tag value.

Note For a list of all conditional expressions available within an **if** statement, see the **if** command.

A tag is a 32-bit integer that can be associated with a given route within the RIB.

The **eq** operator matches either a specific tag value or a parameter value. Its variants **ge** and **le** match a range of tag values that are either greater than or equal to or less than or equal to the supplied value or parameter.

Task ID

Task ID	Operations
route-policy	read, write

Examples

In the following example, if the tag equals 10, then the condition returns true:

```
RP/0/RSP0RP0/CPU0:router(config-rpl)# if tag eq 10 then
```

tag in

To match a tag entry in a named tag set or inline tag set, use the **tag in** command in route-policy configuration mode.

tag in {*tag-set-nameinline-tag-setparameter*}

Syntax Description

tag-set-name	Name of a tag set. The tag-set accepts 32-bit Integer value. Range is 0 to 4294967295.
inline-tag-set	Inline tag set. The inline tag set must be enclosed in parentheses.
parameter 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 4.3.0	This command was introduced.

Usage Guidelines

Use the **tag in** command as a conditional expression within an **if** statement to match a tag entry in a named tag set or inline tag set.

Note For a list of all conditional expressions available within an **if** statement, see the **if** command.

This command takes either a named tag set or an inline tag set value as an argument. The condition returns true if the tag entry matches any entry in the tag set or inline tag set. An attempt to match a tag using a tag set that is defined but contains no elements returns false.

The routing policy language (RPL) provides the ability to test tags for a match to a list of tag match specifications using the **in** operator. The **tag in** command is protocol-independent.

Task ID

Task ID	Operations
route-policy	read, write

Examples

In the following example, a tag set named `my-tag-set` is defined and a route policy named `use-tag-in` is created. Within the `use-tag-in` route policy, the **tag in** command is used within an **if** statement to learn if the tag is in the tag-set named `my-tag-set`. If it is, then local preference is set to 100. If it is not in `my-tag-set` but does match the next tag specifications, then local preference is set to 200.

```
RP/0/RSP0/CPU0:router(config)#tag-set my-tag-set
RP/0/RSP0/CPU0:router(config-tag)#1000
```

```
RP/0/RSP0/CPU0:router(config-tag)#3000
RP/0/RSP0/CPU0:router(config-tag)#end-set

RP/0/RSP0/CPU0:router(config)#route-policy use-tag-in
RP/0/RSP0/CPU0:router(config-rpl)#if tag in my-tag-set then
RP/0/RSP0/CPU0:router(config-rpl-if)#set local-preference 100
RP/0/RSP0/CPU0:router(config-rpl-if)#elseif tag in (2000, 4000) then
RP/0/RSP0/CPU0:router(config-rpl-elseif)#set local-preference 200
RP/0/RSP0/CPU0:router(config-rpl-elseif)#endif
RP/0/RSP0/CPU0:router(config-rpl)#end policy
```

tag-set

To enter tag set configuration mode and define a tag set, use the **tag-set** command in global configuration mode. To remove a named tag set, use the **no** form of this command.

tag-set *name*
no tag-set *name*

Syntax Description

name	Name of a tag set.

Command Default

None

Command Modes

Global configuration

Command History

Release	Modification
Release 4.3.0	This command was introduced.

Usage Guidelines

Use the **tag-set** command to enter tag set configuration mode and define a tag set. A tag-set is a 32-bit integer that can be associated with a given route within the RIB.

Task ID

Task ID	Operations
route-policy	read, write

Examples

In the following example, a tag set named `my-tag-set` is defined and a route policy named `use-tag-in` is created. Within the `use-tag-in` route policy, the **tag in** command is used within an **if** statement to learn if the tag is in the tag-set named `my-tag-set`. If it is, then local preference is set to 100. If it is not in `my-tag-set` but does match the next tag specifications, then local preference is set to 200.

```
RP/0/RSP0/CPU0:router(config)#tag-set my-tag-set
RP/0/RSP0/CPU0:router(config-tag)#1000
RP/0/RSP0/CPU0:router(config-tag)#3000
RP/0/RSP0/CPU0:router(config-tag)#end-set

RP/0/RSP0/CPU0:router(config)#route-policy use-tag-in
RP/0/RSP0/CPU0:router(config-rpl)#if tag in my-tag-set then
RP/0/RSP0/CPU0:router(config-rpl-if)#set local-preference 100
RP/0/RSP0/CPU0:router(config-rpl-if)#elseif tag in (2000, 4000) then
RP/0/RSP0/CPU0:router(config-rpl-elseif)#set local-preference 200
RP/0/RSP0/CPU0:router(config-rpl-elseif)#endif
RP/0/RSP0/CPU0:router(config-rpl)#end policy
```

unsuppress-route

To indicate that a given component of a BGP aggregate should be unsuppressed, use the **unsuppress-route** command in route-policy configuration mode.

unsuppress-route

Syntax Description

This command has no arguments or keywords.

Command Default

No default behavior or values

Command Modes

Route-policy 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **unsuppress-route** command to indicate that a given component of an aggregate should be unsuppressed, that is, allowed to be advertised by BGP again. This command affects routes that have been suppressed in the generation of BGP aggregates. If the request to unsuppress a route is encountered in a policy at a neighbor-out attach point, it guarantees that the routes that it affects are advertised to that neighbor even if that route was suppressed using the **suppress-route** command in a policy at the aggregation attach point.

The **unsuppress-route** command can be used as an action statement within an **if** statement. For a list of all action statements available within an **if** statement, see the **if** command.

Task ID

Task ID	Operations
route-policy	read, write

Examples

In the following example, if the destination is in 10.1.0.0/16, then the route is not advertised:

```
RP/0/RSP0/CPU0:router(config)# route-policy check-aggregate
RP/0/RSP0/CPU0:router(config-rpl)# if destination in (10.1.0.0/16) then
RP/0/RSP0/CPU0:router(config-rpl-if)# unsuppress-route

RP/0/RSP0/CPU0:router(config-rpl-if)# endif
RP/0/RSP0/CPU0:router(config-rpl)# end-policy
```

Assuming that the policy is attached at a neighbor-out attach point, if the route 10.1.0.0/16 was suppressed in a policy at an aggregation attach point, 10.1.0.0/16 is advertised to the neighbor. Routes continue to be suppressed in advertisements to other BGP neighbors unless a specific policy is attached to unsuppress the route.

Related Commands

Command	Description
suppress-route, on page 1889	Indicates that a given component of a BGP aggregate should be suppressed.

var globalVar*N*

To assign a value to route-policy global variable "globalVar1, globalVar2, globalVar3, globalVar4, and globalVar5", use the **var globalVar***N* command in route-policy configuration mode.

var globalVarN *{number | parameter}*

Syntax Description

number	Value assigned to a 32-bit unsigned integer. Range is from 1 to 4294967295.
parameter	Parameter name. The parameter name must be preceded with a "$."

Command Default

If the var globalVar*N* statement is not present then the value for globalVar*N* is zero.

Command Modes

Route-policy configuration

Command History

Release	Modification
Release 5.1.3	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.

The globalVar*N* variables can be used to control policy execution flow. You can assign a weightage to mark the flow and check the value using an if statement.

Note The var globalVar*N* represents statements "var globalVar1" , "var globalVar2" , "var globalVar3" , "var globalVar4" and "var globalVar5". This command can be used as an assignment statement within an **if** statement or in child policies and the value can be used to compare in an **if** statement.

Task ID

Task ID	Operation
route-policy	read, write

Example

In the following example, the globalVar1 is set to 123 and globalVar2 is set to the value stored in parameter $param.

```
RP/0/0/CPU0:ios(config-rpl)#var globalVar1 123
RP/0/0/CPU0:ios(config-rpl)#var globalVar2 $param
```

vpn-distinguisher is

To match a specific Border Gateway Protocol (BGP) VPN distinguisher, use the **vpn-distinguisher is** command in route-policy configuration mode.

vpn-distinguisher is {*numberparameter*}

Syntax Description

number	Value assigned to a 32-bit unsigned integer. Range is from 1 to 4294967295.
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 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **vpn-distinguisher is** command as a conditional expression within an **if** statement to test the value of the origin attribute.

A VPN distinguisher is used in Layer 3 VPN networks for enhanced individual VPN control and to avoid route target mapping at AS boundaries in inter-AS VPN networks. Route target extended communities are removed at neighbor outbound and the VPN distinguisher value is applied on the BGP route as an extended community. When the route is received on a neighboring router in another AS, the VPN distinguisher is removed and mapped to a route target extended community.



Note For a list of all conditional expressions available within an **if** statement, see the **if** command.

This command can be parameterized.

Task ID

Task ID	Operations
route-policy	read, write

Examples

In the following example, the origin is tested within an **if** statement to learn if it is either **igp** or **egp** :

```
RP/0/RSP0/CPU0:router(config-rpl)# if origin is igp or origin is egp then
```

In the following example, a parameter is used to match a specific origin type:

```
RP/0/RSP0/CPU0:router(config)# route-policy bar($origin)
RP/0/RSP0/CPU0:router(config-rpl)# if origin is $origin then
RP/0/RSP0/CPU0:router(config-rpl-if)# set med 20
RP/0/RSP0/CPU0:router(config-rpl-if)# endif
RP/0/RSP0/CPU0:router(config-rpl)#
```



Static Routing Commands

This module describes the commands used to establish static routes on Cisco ASR 9000 Series Aggregation Services Routers .

For detailed information about static routing concepts, configuration tasks, and examples, see the *Implementing Static Routes on* Cisco ASR 9000 Series Router module in the *Routing Configuration Guide for Cisco ASR 9000 Series Routers*.

- address-family (static), on page 1900
- maximum path (static), on page 1902
- metric (static), on page 1904
- route (static), on page 1906
- router static, on page 1910
- track (static), on page 1912
- vrf (static), on page 1914

address-family (static)

To enter various address family configuration modes while configuring static routes, use the **address-family** command in the appropriate configuration mode. To disable support for an address family, use the **no** form of this command.

address-family {**ipv4** | **ipv6**} {**unicast** | **multicast**}
no address-family {**ipv4** | **ipv6**} {**unicast** | **multicast**}

Syntax Description

ipv4	Specifies IP Version 4 address prefixes.
ipv6	Specifies IP Version 6 address prefixes. This option is available only in static router configuration mode.
unicast	Specifies unicast address prefixes.
multicast	Specifies multicast address prefixes. This option is available only in static router configuration mode.

Command Default

All static routes belong to the default VRF if you enter address family configuration mode without entering VRF configuration mode.

Command Modes

Router static configuration

VRF router static 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **address-family** command to enter various address family configuration modes while configuring static routing sessions. From address family configuration mode, you can configure static routes using the **route** command.

Task ID

Task ID	Operations
static	read, write

Examples

The following example shows how to enter IPv6 unicast address family mode:

```
RP/0/RSP0/CPU0:router(config)# router static
RP/0/RSP0/CPU0:router(config-static)# address-family ipv6 unicast
```

```
RP/0/RSP0/CPU0:router(config-static-afi)#
```

Related Commands

Command	Description
route (static), on page 1906	Establishes a static route.

maximum path (static)

To change the maximum number of allowable static routes, use the **maximum path** command in static router configuration mode. To remove the **maximum path** command from the configuration file and restore the system to its default condition, use the **no** form of this command.

maximum path {**ipv4** | **ipv6**} *value*
no maximum path {**ipv4** | **ipv6**} *value*

Syntax Description

ipv4 | **ipv6**	Specifies IP Version 4 (IPv4) or IP Version 6 (IPv6) address prefixes.
value	Maximum number of static routes for the given AFI. The range is 1 to 140000.

Command Default

value : 4000

Command Modes

Static router 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

If you use the **maximum path** command to reduce the configured maximum allowed number of static routes for a given table below the number of static routes currently configured, the change is rejected. In addition, if you commit a batch of routes that would, when grouped, push the number of static routes configured above the maximum allowed, the first *n* routes in the batch and the number previously configured are accepted, and the remainder rejected. The *n* argument is the difference between the maximum number allowed and the number previously configured.

Task ID

Task ID	Operations
static	read, write

Examples

The following example shows how to set the maximum number of static IPv4 routes to 100000:

```
RP/0/RSP0/CPU0:router(config-static)# maximum path ipv4 100000
```

The following example shows how to remove the preceding configuration and set the maximum number of static IPv4 routes back to the default:

```
RP/0/RSP0/CPU0:router(config-static)# no maximum path ipv4 100000
```

Related Commands

Command	Description
route (static), on page 1906	Enters static router configuration mode.
show route static	Displays the static routes in a routing table.

metric (static)

To set metric values for static routes, use the **metric** command on a route after you have entered address family configuration mode. To disable metric values, use the **no** form of this command.

ip-address | *interface-path-id* [**metric** *metric-value*]
no *ip-address* | *interface-path-id* [**metric** *metric-value*]

Syntax Description

ip-address	The IP address is required and, if the interface-path-id arguments are not specified, then the metric option is not displayed.
interface-path-id	Physical interface or virtual interface path-id.
metric *metric-value*	Set a metric value for a static route based on the available bandwidth. Ideal range is 0 to 32.

Command Default

Metric values are not set.

Command Modes

Static IPv4 address family

Static IPv6 address family

Static VRF IPv4 address family

Static VRF IPv6 address family

Command History

Release	Modification
Release 6.0.1	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.

Use the **metric** command to provide metric values for a static route with weights.

Task ID

Task ID	Operation
static	read, write

Example

This example shows how to set the metric value as 5 for a route:

```
RP/0/RSP0/CPU0:router(config)# router static
RP/0/RSP0/CPU0:router(config-static)# address-family ipv4 unicast
```

```
RP/0/RSP0/CPU0:router(config-static-afi)# 1.1.1.1/32 gigabitEthernet 0/0/0/1 metric 5
```

route (static)

To establish static routes, use the **route** command in address family configuration mode. To remove the **route** command from the configuration, use the **no** form of this command.

prefix/mask [**vrf** *vrf-name*] {*ip-address* | *type* *interface-path-id* [*ip-address* | *type* *interface-path-id*] [**behavior** | *utef*] [**track** *track-object-name*] [**tunnel-id** *tunnel-id*] [**vrflabel** *vrf-label*] [*distance*] [**description** *text*] [**tag** *tag*] [**permanent**]}

Syntax Description

prefix / *mask*	IP route prefix and prefix mask for the destination. The network mask can be specified in either of two ways: • The network mask can be a four-part, dotted-decimal address. For example, 255.0.0.0 indicates that each bit equal to 1 means the corresponding address bit is a network address. • The network mask can be indicated as a slash (/) and number. For example, /8 indicates that the first 8 bits of the mask are 1s, and the corresponding bits of the address are the network address.

vrf *vrf-name*	(Optional) Specifies a destination VRF. This option is available when IPv4 address families are specified.
	The following names cannot be used: all, default, and global.
	The following example shows how to configure IPv4 VRF:
	`router static`
	` address-family ipv4 unicast`
	`  10.1.1.0/24 vrf vrf_a 192.168.1.1`
	`router static`
	` vrf vrf_a`
	`  address-family ipv4 unicast`
	`   10.1.1.0/24 192.168.1.1`

ip-address	IP address of the next hop that can be used to reach that network. • For IPv4 address–the IP address is required, not optional, if the interface type and interface-path-id arguments are not specified. You can specify an IP address and an interface type and interface path. • For IPv6 link-local address–the interface type and interface-path-id arguments are required. The route is not valid, if the interface type and interface-path-id arguments are not specified. **Note** A forwarding router's IP address or an interface or virtual interface path ID can be configured, in any order.
behavior	Define the behavior by selecting the appropriate option.
utef	Static programs the utef SID with the context of controller address.

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. **Note** A forwarding router's IP address or an interface or virtual interface path ID can be configured, in any order.
distance	(Optional) Administrative distance. Range is 1 to 254.
description *text*	(Optional) Specifies a description of the static route.
tag *tag*	(Optional) Specifies a tag value that can be used as a match for controlling redistribution using route policies. Range is 1 to 4294967295.
permanent	(Optional) Specifies that the route is not removed from the routing table, even if the next-hop interface shuts down or next-hop IP address is not reachable.
track *track-object-name*	Enables object tracking for static route.
tunnel-id *tunnel-id*	Specifies a Tunnel ID.
vrflabel *vrf-label*	Specifies a VRF label.

Command Default

No static route is established.

vrf *vrf-name* : If no VRF is specified, the vrf where the configuration takes place is used.

Command Modes

Address family configuration

Command History

Release	Modification
Release 24.1.1	The *utef* keyword was introduced.
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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

A static route is appropriate when the software cannot build a route to the destination, dynamically.

Static routes have a default administrative distance of 1, in which a low number indicates a preferred route. By default, static routes are preferred to routes learned by routing protocols. You can configure an administrative distance with a static route if you want the static route to be overridden by dynamic routes. For example, you

could have routes installed by the Open Shortest Path First (OSPF) protocol with an administrative distance of 120. To have a static route that would be overridden by an OSPF dynamic route, specify an administrative distance greater than 120.

The routing table considers the static routes that point to an interface as "directly connected." Directly connected networks are advertised by IGP routing protocols if a corresponding **interface** command is contained under the router configuration stanza of that protocol.

A static route is always associated with a VPN routing and forwarding (VRF) instance. The VRF can be the default VRF or a specified VRF. Specifying a VRF allows you to enter VRF configuration mode where you can configure a static route. If you do not specify a VRF you can configure a default VRF static route.

Use the **router static** command to configure static routes. To configure a static route, you must enter router static configuration mode and then enter an address family configuration mode or VRF configuration mode. See the **vrf (static)** command for information on configuring a static route in VRF configuration mode. After you enter an address family mode, you can enter multiple static routes. The following example shows how to configure multiple static routes in IPv4 and IPv6 address family configuration modes:



Note You cannot create a VRF named default, but you can reference the default VRF.

Task ID

Task ID	Operations
static	read, write

The following example shows how to configure the Sink node:

```
Router(config)#router static
Router(config-static)# address-family ipv6 unicast
Router(config-static-afi)#fccc:cc00:9000:fef1::/64 segment-routing srv6 endpoint behavior
utef controller-address fccc:cc00:7::
```

Examples

The following example shows how to configure IPv6 unicast address family static routes:

```
Router(config)# router static
Router(config-static)# address-family ipv6 unicast
Router(config-static-afi)# 2b11::327a:7b00/120 GigabitEthernet0/2/0/7
Router(config-static-afi)# 2b11::327a:7b00/120 GigabitEthernet0/6/0/0
Router(config-static-afi)# 2b11::327a:7b00/120 2b11::2f01:4c
Router(config-static-afi)# 2b11::327a:7b00/120 2b11::2f01:4d
Router(config-static-afi)# 2b11::327a:7b00/120 2b11::2f01:4e
Router(config-static-afi)# 2b11::327a:7b00/120 2b11::2f01:4f
Router(config-static-afi)# 2b11::327a:7b00/120 2b11::2f01:50
```

Related Commands

Command	Description
address-family (static), on page 1900	Enters address family configuration mode.

Command	Description
network (BGP)	Specifies a list of networks for the BGP routing process.
show route	Displays the current contents of the routing table.
show route static	Displays the static routes in a routing table.
show route summary	Displays the current contents of the routing table in summary format.
vrf (static)	Enters VRF static route configuration mode.

router static

To enter static router configuration mode, use the **router static** command in global configuration mode. To remove all static route configurations and terminate the static routing process, use the **no** form of this command.

router static
no router static

Syntax Description

This command has no arguments or keywords.

Command Default

No static routing process is enabled.

Command Modes

Global 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 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
static	read, write
bgp, ospf, hsrp, isis, vrrp, multicast, or network	read, write

Examples

The following example shows how to enter static router configuration mode:

```
RP/0/RSP0/CPU0:router(config)# router static
RP/0/RSP0/CPU0:router(config-static)#
```

Related Commands

Command	Description
address-family (static), on page 1900	Enters address family configuration mode.
show route	Displays the current contents of the routing table.
show route static	Displays the static routes in a routing table.
show route summary	Displays the current contents of the routing table in summary format.
route (static), on page 1906	Establishes a static route.

Command	Description
vrf (static)	Enters VRF static route configuration mode.

track (static)

To enable object racking for static routes, use the **track** command in an appropriate configuration mode. To disable object tracking, use the **no** form of this command.

ip-address / length [**vrf** *vrf-name*] {*nexthop* | *type interface-path-id*} [**track** *object-name*]
no *ip-address / length* [**vrf** *vrf-name*] {*nexthop* | *type interface-path-id*} [**track** *object-name*]

Syntax Description

ip-address / length	Specifies the destination prefix.
vrf *vrf-name*	Specifies vrf for the nexthop.. This option is available when IPv4 and IPv6 address families are specified. These names cannot be used for *vrf-name*: all, default, and global.
nexthop	Specifies the forwarding router's address.
type	Specifies the 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.
object-name	Specifies the name of the object to be tracked. Up to 32 characters are allowed for specifying an object name.

Command Default

Object tracking is disabled. No object is specified for tracking.

Command Modes

Static IPv4 address family

Static VRF IPv4 address family

Command History

Release	Modification
Release 4.2.1	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.

If the specified objects are not existing in IOS XR database, those objects are forward-referenced until created. However, objects that are forward-referenced by the tracking configuration are treated as 'down'.

When a failure between the first hop router and core occurs, the tracking object will go Down. The static router withdraws the primary paths, which were tracking this object and installs the back-up paths, if present.

Task ID

Task ID	Operation
static	read, write

This example shows how to enable tracking for an object *object1*:

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#router static
RP/0/RSP0/CPU0:router(config-static)#address-family ipv4 unicast
RP/0/RSP0/CPU0:router(config-static-afi)#100.0.24.0/24 204.0.23.2 track object1
```

vrf (static)

To configure a VPN routing and forwarding (VRF) instance and enter VRF configuration mode, use the **vrf** command in router configuration 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.

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

Static router 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 appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

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.

A static route is always associated with a VRF, which is entirely user configurable. Static route is unique within a VRF. A static route can point to a next-hop interface, next-hop IP address, or both, which can be resided in the same VRF configured for the static route or in a different VRF. For example, routes 172.168.40.0/24 and 172.168.50.0/24 are configured as follows:

```
router static
vrf vrf_A
address ipv4 unicast
172.168.40.0/24 loopback 1
172.168.50.0/24 vrf vrf_B 192.168.1.2
```

Routes 172.168.40.0/24 and 172.168.50.0/24 belong to vrf_A. Route 172.168.50.0/24 is not installed in vrf_A until next-hop 192.168.1.2 (a vrf_B route) is reachable.

If you are configuring a default VRF route, you do not need to enter VRF configuration mode. For example, routes 192.168.1.0/24 and 192.168.2.0/24 are configured as follows:

```
router static
 address ipv4 unicast
  192.168.1.0/24 loopback 5
  192.168.2.0/24 10.1.1.1
```

Routes 192.168.1.0/24 and 192.168.2.0/24 are default VRF routes.

Note You cannot create a VRF named default, but you can reference the default VRF.

You must remove IPv4/IPv6 addresses from an interface prior to assigning, removing, or changing a VRF on an IP interface. If this is not done in advance, any attempt to change the VRF on an IP interface is rejected.

Task ID

Task ID	Operations
static	read, write

Examples

The following example shows how to configure a VRF instance and enter VRF configuration mode:

```
RP/0/RSP0/CPU0:router(config)# router static
RP/0/RSP0/CPU0:router(config-static)# vrf vrf-1
RP/0/RSP0/CPU0:router(config-static-vrf)#
```

Related Commands

Command	Description
address-family (static), on page 1900	Enters address family configuration mode and allows you to configure a static route.



RCMD Commands

This module describes the commands used to configure and diagnose RCMD.

For detailed information about RCMD concepts, configuration tasks, and examples, see the *Implementing RCMD* module in the *Routing Configuration Guide for Cisco ASR 9000 Series Routers*.

- router-convergence, on page 1918
- monitor-convergence (IS-IS), on page 1919
- monitor-convergence (OSPF), on page 1920
- collect-diagnostics (RCMD), on page 1921
- event-buffer-size (RCMD), on page 1923
- max-events-stored (RCMD), on page 1924
- monitoring-interval (RCMD), on page 1925
- node disable (RCMD), on page 1927
- prefix-list (monitor-convergence IS-IS), on page 1929
- prefix-list (monitor-convergence OSPF), on page 1931
- priority (RCMD), on page 1933
- protocol (RCMD), on page 1935
- show rcmd isis event prefix, on page 1936
- show rcmd ospf event prefix, on page 1938
- show rcmd ospf event spf , on page 1940
- storage-location, on page 1943
- track-external-routes, on page 1945
- track-summary-routes, on page 1946

router-convergence

To configure route convergence monitoring and enter router convergence monitoring and diagnostics (rcmd) configuration mode, use the **router-convergence** command in global configuration mode. To remove all router convergence monitoring configurations and exit the rcmd mode, use the **no** form of this command.

router-convergence [**disable**]
no router-convergence

Syntax Description

disable	[Optional] Disables the monitoring of route convergence on the entire router.

Command Default

RCMD is disabled.

Command Modes

Global configuration

Command History

Release	Modification
Release 4.2.0	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
rcmd	read, write

This example shows how to configure router-convergence command and enable rcmd configuration mode:

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#router-convergence
RP/0/RSP0/CPU0:router(config-rcmd)#
```

Related Commands

Command	Description
monitor-convergence (IS-IS), on page 1919	Enables route convergence monitoring for IS-IS protocol.
monitor-convergence (OSPF), on page 1920	Enables OSPF route convergence monitoring.

monitor-convergence (IS-IS)

To enable route convergence monitoring for IS-IS protocol, use the **monitor-convergence** command in address family configuration mode. To disable, route convergence monitoring, use the **no** form of this command.

monitor-convergence
no monitor-convergence

Syntax Description

This command has no keywords or arguments.

Command Default

Route convergence monitoring is disabled.

Command Modes

Address family IPv4 unicast

Address family IPv6 unicast

Address family IPv4 multicast

Address family IPv6 multicast

Command History

Release	Modification
Release 4.2.0	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
isis	read, write

This example shows how to configure route convergence monitoring for IS-IS under IPv6 unicast SAFI:

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#router isis isp
RP/0/RSP0/CPU0:router(config-isis)#address-family ipv6 unicast
RP/0/RSP0/CPU0:router(config-isis-af)#monitor-convergence
```

Related Commands

Command	Description
router-convergence, on page 1918	Configures route convergence monitoring and enters router convergence monitoring and diagnostics (rcmd) configuration mode.
monitor-convergence (OSPF), on page 1920	Enables OSPF route convergence monitoring.

monitor-convergence (OSPF)

To enable OSPF route convergence monitoring, use the **monitor-convergence** command in router OSPF configuration mode. To disable OSPF route convergence monitoring, use the **no** form of this command.

monitor-convergence
no monitor-convergence

Syntax Description

This command has no keywords or arguments.

Command Default

Monitor Convergence is disabled.

Command Modes

Router configuration

Command History

Release	Modification
Release 4.2.0	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
ospf	read, write

This example shows how to enable route convergence monitoring for an OSPF process:

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#router ospf 100
RP/0/RSP0/CPU0:router(config-ospf)#monitor-convergence
```

Related Commands

Command	Description
router-convergence, on page 1918	Configures route convergence monitoring and enters router convergence monitoring and diagnostics (rcmd) configuration mode.
monitor-convergence (IS-IS), on page 1919	Enables route convergence monitoring for IS-IS protocol.

collect-diagnostics (RCMD)

To collect diagnostics on specified node, use the **collect-diagnostic** command in router-convergence configuration mode. To disable collection of diagnostics, use the **no** form of this command.

collect-diagnostics *location*
no collect-diagnostics *location*

Syntax Description

location	Specifies the line-card location.

Command Default

Diagnostics collection is disabled.

Command Modes

Router-convergence configuration

Command History

Release	Modification
Release 4.2.0	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.

For enabling diagnostics collection on specific line-card locations, you can user can specify partially qualified semantics. However, it is not allowed to configure over-lapping locations so as to avoid errors. The following combinations of Rack and Slot are accepted:

- */*/*
- R/*/*
- R/S/*

If a wildcard combination for any location is already disabled, then any other combination that overlaps with it would be rejected. For example,

- If */*/* is disabled, then all other disable commands will be rejected
- If R/*/* is disabled, then disable for */*/* and R/S/* will be rejected
- If R/S/* is disabled, then disable for */*/* and R/*/* will be rejected

Task ID

Task ID	Operation
rcmd	read, write

This example shows how to enable RCMD diagnostics collection on node 0/3/CPU0:

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router#router-convergence
RP/0/RSP0/CPU0:router(config-rcmd)#collect-diagnostics 0/3/CPU0
```

Related Commands

Command	Description
router-convergence, on page 1918	Configures route convergence monitoring and enters router convergence monitoring and diagnostics (rcmd) configuration mode.

event-buffer-size (RCMD)

To specify event buffer size (in terms of number of events) for storing event traces, use the **event-buffer-size** command in router-convergence configuration mode. To disable buffer size configuration, use the **no** form of this command.

event-buffer-size *number*
no event-buffer-size

Syntax Description

number	Specifies the Specify the number of events. The range is 100 to 500.

Command Default

100 events.

Command Modes

Router-convergence configuration

Command History

Release	Modification
Release 4.2.0	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.

The event-buffer-size configuration controls the ltrace buffer size. Ltraces will be stored for only the configured number of events. The default is 100 events and can be set based on the expected churn in the network. Value for event buffer impact memory usage on all RPs and monitored LCs.

Task ID

Task ID	Operations
rcmd	read, write

Examples

This example shows how to configure event buffer size as 500 events:

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#router-convergence
RP/0/RSP0/CPU0:router(config-rcmd)#event-buffer-size 500
```

Related Commands

Command	Description
router-convergence, on page 1918	Configures route convergence monitoring and enters router convergence monitoring and diagnostics (rcmd) configuration mode.

max-events-stored (RCMD)

To configure maximum number of events to be stored in the RCMD server, use the **max-events-stored** command in router-convergence configuration mode. To remove the number of events to be stored, use the **no** form of this command.

max-events-stored *number*

Syntax Description

number	Specifies the maximum number of events stored. The range is 10 to 500.

Command Default

100 events.

Command Modes

Router-convergence configuration

Command History

Release	Modification
Release 4.2.0	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.

The max-events-stored configuration controls the number of events that are stored in RCMD server, before the older events are deleted. The default is 100 events and can be set based on the expected churn in the network. Value for events stored impact memory usage by RCMD server.

Task ID

Task ID	Operations
rcmd	read, write

Examples

This example shows how to configure 500 number of events to be stored in RCMD server:

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#router-convergence
RP/0/RSP0/CPU0:router(config-rcmd)#max-events-stored 500
```

Related Commands

Command	Description
router-convergence, on page 1918	Configures route convergence monitoring and enters router convergence monitoring and diagnostics (rcmd) configuration mode.

monitoring-interval (RCMD)

To configure interval (in minutes) in which to collect logs, use the **monitoring-interval** command in router-convergence configuration mode. To disable monitoring interval configuration, use the **no** form of this command.

monitoring-interval *minutes*
no monitoring-interval *minutes*

Syntax Description

minutes	Specifies the interval (in minutes) for collecting logs. The range is 5 to 120 minutes.

Command Default

Periodic monitoring interval is 15 minutes.

Command Modes

Router-convergence configuration

Command History

Release	Modification
Release 4.2.0	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.

The monitoring-interval timer controls the collection, processing, and archival (optional) of convergence data by RCMD server.

Periodic processing can get triggered if number of events detected exceed configured sizing parameters to prevent loss of data . However, this is not guaranteed since the mechanism is throttled.

To collect logs manually, use the **rcmd trigger-data-collect** command. Syslogs are generated when high churn is detected and collection mechanism is getting throttled. This indicates possible loss of data for some events. Throttling mechanism is for one processing every minute.

Task ID

Task ID	Operations
rcmd	read, write

Examples

This example shows how to configure monitoring interval as 5 minutes:

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#router-convergence
RP/0/RSP0/CPU0:router(config-rcmd)#monitoring-interval 5
```

Related Commands

Command	Description
router-convergence, on page 1918	Configures route convergence monitoring and enters router convergence monitoring and diagnostics (rcmd) configuration mode.

node disable (RCMD)

To disable monitoring of route convergence on specified location, use the **node disable** command in router-convergence configuration mode. To reinstate, monitoring on specified location, use the **no** form of this command.

node *node-id* **disable**
no node *node-id* **disable**

Syntax Description

node-id	Specifies line card locations for which RCMD monitoring be disabled. Disables RCMD monitoring on the specified node. No data from this node will be available in the reports that are generated. You can enter specific LCs or use wild cards.

Command Default

Update times are gathered and reported for all LCs. Diagnostic mode is disabled on all LCs.

Command Modes

Router-convergence configuration

Command History

Release	Modification
Release 4.2.0	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.

Disable monitoring on specific LCs or racks for better scalability. Disable monitoring on LCs whose update times is not going to impact the core IGP/LDP convergence that RCMD is measuring.

On LCs where monitoring is enabled, the diagnostic mode can be enabled (with threshold value) for triggering script using EEM infra for debug data collection from the router. Use diagnostic mode only for debugging purpose since it is more CPU intensive as compared to normal RCMD monitoring.

Only the following combinations of Rack and Slot are acceptable:

- */*/*
- R/*/*
- R/S/*

If a wildcard combination for any location is already disabled, then any other combination that overlaps with it would be rejected. For example,

- If */*/* is disabled, then all other disable commands would be rejected
- If R/*/* is disabled, then disable for */*/* and R/S/* would be rejected
- If R/S/* is disabled, then disable for */*/* and R/*/* would be rejected

Task ID

Task ID	Operations
rcmd	read, write

Examples

This example shows how to disable monitoring on all nodes with Rack 0 and any slot (used wild card *) :

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#router-convergence
RP/0/RSP0/CPU0:router(config-rcmd)#node 0/*/* disable
```

Related Commands

Command	Description
router-convergence, on page 1918	Configures route convergence monitoring and enters router convergence monitoring and diagnostics (rcmd) configuration mode.

prefix-list (monitor-convergence IS-IS)

To enable individual prefix monitoring for IS-IS prefixes, use the prefix-list command in Router IS-IS monitor-convergence configuration mode. To disable individual prefix monitoring, use the no form of this command.

prefix-list *prefix-list-name*
no prefix-list *prefix-list-name*

Syntax Description

prefix-list-name	Specifies the name of an IS-IS prefix-list. **Note** Configure a prefix-list under IPv4 or IPv6 using the prefix-list (IP Addresses) command to use for prefix monitoring.

Command Default

All IS-IS prefixes are marked for monitoring, if the prefix-list is not configured

Command Modes

Router IS-IS monitor-convergence

Command History

Release	Modification
Release 4.3.0	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.

To enable monitoring of individual prefixes, first configure a prefix-list using the {ipv4 | ipv6} prefix-list command. Then, use this prefix list with the prefix-list (monitor-convergence IS-IS).

Task ID

Task ID	Operation
isis	read, write

This example shows how to enable IS-IS prefix monitoring:

First, configure a prefix-list:

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#ipv4 prefix-list isis_monitor
RP/0/RSP0/CPU0:router(config-ipv4_pfx)#10 permit 35.0.0.0/8 eq 32
RP/0/RSP0/CPU0:router(config-ipv4_pfx)#commit
RP/0/RSP0/CPU0:router(config-ipv4_pfx)#exit
```

Then, configure the prefix list command under Router IS-IS monitor-convergence configuration mode:

```
RP/0/RSP0/CPU0:router(config)#router isis isp
RP/0/RSP0/CPU0:router(config-isis)#address-family ipv4 unicast
RP/0/RSP0/CPU0:router(config-isis-af)#monitor-convergence
RP/0/RSP0/CPU0:router(config-isis-af-rcmd)#prefix-list isis_monitor
```

prefix-list (monitor-convergence OSPF)

To enable individual prefix monitoring for OSPF prefixes, use the **prefix-list** command in Router OSPF monitor-convergence configuration mode. To disable individual prefix monitoring, use the **no** form of this command.

prefix-list *prefix-list-name*
no prefix-list *prefix-list-name*

Syntax Description

prefix-list-name	Specifies the name of an OSPF prefix-list. **Note** Configure a prefix-list under IPv4 or IPv6 using the prefix-list (IP Addresses) command to use for prefix monitoring.

Command Default

All OSPF prefixes are marked for monitoring, if the prefix-list is not configured.

Command Modes

Router OSPF monitor-convergence

Command History

Release	Modification
Release 4.3.0	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.

To enable monitoring of individual prefixes, first configure a prefix-list using the {ipv4 | ipv6} prefix-list command. Then, use this prefix list with the prefix-list (monitor-convergence OSPF).

Task ID

Task ID	Operation
ospf	read, write

This example shows how to enable OSPF prefix monitoring:

First, configure a prefix-list:

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#ipv4 prefix-list ospf_monitor
RP/0/RSP0/CPU0:router(config-ipv4_pfx)#10 permit 35.0.0.0/8 eq 32
RP/0/RSP0/CPU0:router(config-ipv4_pfx)#commit
RP/0/RSP0/CPU0:router(config-ipv4_pfx)#exit
```

Then, configure the prefix list command under Router OSPF monitor-convergence configuration mode:

```
RP/0/RSP0/CPU0:router(config)#router ospf 1
RP/0/RSP0/CPU0:router(config-ospf)#monitor-convergence
RP/0/RSP0/CPU0:router(config-ospf-af-rcmd)#prefix-list ospf_monitor
```

priority (RCMD)

To configure RCMD reporting parameters for low/high/critical/medium priority updates, use the **priority** command in RCMD protocol configuration mode. To disable setting up priority use the **no** form of this command.

priority {**Critical** | **High** | **Low** | **Medium**} [**disable**] [**leaf-network** *leaf-network-number*] [**threshold** *value*]
no priority {**Critical** | **High** | **Low** | **Medium**}

Syntax Description

Critical	Configures the monitoring of route convergence for critical routes.
High	Configures the monitoring of route convergence for high priority routes.
Low	Configures the monitoring of route convergence for low priority routes.
Medium	Configures the monitoring of route convergence for medium priority routes.
disable	Disables the monitoring of route convergence for specified priority.
leaf-network	Configures the monitoring of route convergence for leaf networks. Lists up to 100 leaf networks that were added or deleted as part of SPF.
leaf-network-number	Specifies the maximum number of leaf networks monitored. The range is 10 to 100.
threshold	Sets the threshold value for convergence in milliseconds. If the convergence time exceeds this configured value, diagnostics collection will be triggered.
value	Specifies the threshold value (in msec). The range is 0 to 4294967295.

Command Default

None

Command Modes

Router-convergence protocol configuration

Command History

Release	Modification
Release 4.2.0	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.

Use the **priority** command for collecting data and applying threshold for particular protocol and prefix priority.

Maximum of 100 leaf networks can get logged because of scalability reasons. No default value available for threshold and this needs to be determined with deployment experience for specific network. Threshold specification is required for triggering diagnostics collection. Disable monitoring for medium and or low priority routes to help scale better. No specific order is guaranteed for leaf networks and first N prefixes that change are logged.

Task ID

Task ID	Operations
rcmd	read, write

Examples

This example shows how Configures the monitoring of route convergence for critical routes for 100 leaf networks and at a threshold value of 1 millisecond for OSPF protocol:

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#router-convergence
RP/0/RSP0/CPU0:router(config-rcmd)#protocol OSPF
RP/0/RSP0/CPU0:router(config-rcmd-proto)#priority high
RP/0/RSP0/CPU0:router(config-rcmd-proto-prio)#leaf-network 100
RP/0/RSP0/CPU0:router(config-rcmd-proto-prio)#threshold 1
```

Related Commands

Command	Description
router-convergence, on page 1918	Configures route convergence monitoring and enters router convergence monitoring and diagnostics (rcmd) configuration mode.

protocol (RCMD)

To specify the protocol for which to configure RCMD parameters, use the **protocol** command in router-convergence configuration mode. To remove the protocol from RCMD, use the **no** form of this command.

protocol {**ISIS** | **OSPF**}
no protocol {**ISIS** | **OSPF**}

Syntax Description

ISIS	Configures parameters related to OSPF protocol within RCMD
OSPF	Configures parameters related to IS-IS protocol within RCMD

Command Default

None

Command Modes

Router-convergence configuration

Command History

Release	Modification
Release 4.2.0	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.

RCMD monitoring needs to be enabled for a specific OSPF or ISIS protocol instance.

Task ID

Task ID	Operations
rcmd	read, write

Examples

This example shows how to enable RCMD parameters for OSPF protocol:

```
RP/0/RSP0/CPU0:router(config)#router
RP/0/RSP0/CPU0:router(config)#router-convergence
RP/0/RSP0/CPU0:router(config-rcmd)#protocol OSPF
RP/0/RSP0/CPU0:router(config-rcmd-proto)#priority high
RP/0/RSP0/CPU0:router(config-rcmd-proto-prio)#leaf-network 100
RP/0/RSP0/CPU0:router(config-rcmd-proto-prio)#threshold 1
```

Related Commands

Command	Description
router-convergence, on page 1918	Configures route convergence monitoring and enters router convergence monitoring and diagnostics (rcmd) configuration mode.

show rcmd isis event prefix

To display the details of the individual IS-IS prefix events, use the show rcmd isis event prefix command in EXEC mode.

show rcmd isis *isis-instance* **event prefix** [*event-numberprefix* | **after** *event_number* | **last** *event_number* | **priority** {**critical** | **high** | **low** | **medium**} | **threshold-exceeded**] [**detail**] [**xml**]

Syntax Description

isis-instance	Specifies the name of an IS-IS instance.
event-number	(Optional) Specifies the number of a specific event that is run. Range is 0-4294967295.
prefix	(Optional) Specifies events with a prefix. Specify prefix in *ip-address/length* format.
after	(Optional) Specifies events after a specific event number.
last	(Optional) Specifies the last number of events. Range is 1-500.
priority	(Optional) Specifies to filter events by priority. • **critical**—Event that has critical priority prefixes. • **high**—Event that has high priority prefixes. • **low**—Event that has low priority prefixes. • **medium**—Event that has medium priority prefixex.
threshold-exceeded	(Optional) Specifies events that have exceeded the threshold.
detail	(Optional) Provides detailed output data.
xml	(Optional) Provides output in XML format

Command Default

None

Command Modes

EXEC

Command History

Release	Modification
Release 4.3.0	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
isis	read

This is sample output from the **show rcmd isis event prefix** command:

```
RP/0/RSP0/CPU0:router#show rcmd isis isp event prefix
```

show rcmd ospf event prefix

show rcmd ospf *ospf-instance* **event prefix** [*event_numberprefix* | **after** *event_number* | **last** *event_number* | **priority** {**critical** | **high** | **low** | **medium**} | **threshold-exceeded**] [**detail**] [**xml**]

Syntax Description

ospf-instance	Specifies the name of an OSPF instance.
event-number	(Optional) Specifies the number of a specific event that is run. Range is 0-4294967295.
prefix	(Optional) Specifies events with a prefix. Specify prefix in *ip-address/length* format.
after	(Optional) Specifies events after a specific event number.
last	(Optional) Specifies the last number of events. Range is 1-500.
priority	(Optional) Specifies to filter events by priority. • **critical**—Event that has critical priority prefixes. • **high**—Event that has high priority prefixes. • **low**—Event that has low priority prefixes. • **medium**—Event that has medium priority prefixex.
threshold-exceeded	(Optional) Specifies events that have exceeded the threshold.
detail	(Optional) Provides detailed output data.
xml	(Optional) Provides output in XML format

Command Default

None

Command Modes

EXEC

Command History

Release	Modification
Release 4.3.0	This command was introduced.

Usage Guidelines

Task ID

Task ID	Operation
ospf	read, write

This is sample output from **show rcmd ospf event prefix** command:

```
OSPF process: 1
```

```
Event: 1

    Prefix: 255.255.255.255/32    Cost: 10              Priority: High
    SPF Event No: 0               Route-Type: Intra     Change-Type: Add

    Nexthop: GigabitEthernet-2/0/0/1 Neighbor: 2.2.2.2  Change-Type: Add
             GigabitEthernet-2/0/0/2 Neighbor: 1.1.1.1  Change-Type: Delete

    Start time: Jan 1 05:32:22.118
    Timeline:
        IP Route Program Time:        Min: 40(0/2/CPU0)    Max: 66(0/1/CPU0)
        MPLS Label Program Time:      Min: 173(0/1/CPU0)   Max: 197(0/3/CPU0)
        Details:
              RIBv4-Enter               6     <offset from Start time>
              RIBv4-Exit               12
              RIBv4-Redist              8
              LDP Enter                10
              LDP Exit                 16
              LSD Enter                27
              LSD Exit                 42
              LC Details(IP Path):
                S 0/1/CPU0             66
                F 0/2/CPU0             40
                  0/3/CPU0             56
              LC Details(MPLS Path):
                F 0/1/CPU0            173
                  0/2/CPU0            174
                S 0/3/CPU0            197
```

show rcmd ospf event spf

To display route convergence monitoring and diagnostics information for OSPF shortest path first events, use the **show rcmd ospf event spf** command in EXEC mode.

show rcmd ospf *ospf-instance* **event spf** [*spf-run* | **after** | **last** | **no-route-change** | **pending** | **route-change** | **threshold-exceeded**] [**detail**] [**xml**]

Syntax Description

ospf-instance	Specifies the OSPF instance number.
spf-run	(Optional) Specifies a specific OSPF SPF run. Range is 0-4294967295.
after	(Optional) Specifies events after a specific number of events. Range is 0-4294967295.
last	(Optional) Specifies the last "N" events. Range for "N" is 1-500.
no-route-change	(Optional) Displays information about events that have no-route-changes.
pending	(Optional) Displays events that are pending for post processing.
route-change	(Optional) Displays events that have route-change.
threshold-exceed	(Optional Displays that have exceeded the threshold.
detail	Optional) Displays detailed information about the SPF event.
xml	(Optional) Displays information in XML format.

Command Default

None

Command Modes

EXEC

Command History

Release	Modification
Release 4.3.0	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

This is sample output from the show rcmd ospf event spf command:

```
RP/0/RSP0/CPU0:router#show rcmd ospf 1 event spf last 1 detail

Event Status:
^  no route change #  threshold exceeded ~ incomplete data * collection pending


OSPF process: 1

SPF run: 16
    Trigger: Apr 10 23:07:01.614  Start: 0            Duration: 2
    Dijkstra Statistics:          Runs: 1             LSA changes: 0
    IA/Ext Statistics:            Runs: 4             LSA processed: 36
    Timeline Summary:
        Priority: Critical
            Route Count:                 Added: 18          Deleted: 0          Modified:
 0
            FRR Coverage:                Routes: 9/9(100%)             Paths: 18/18(100%)

            IP Route Program Time:       Min: 9(0/2/CPU0)              Max: 11(0/3/CPU0)

            MPLS Label Program Time:     Min: 18(0/1/CPU0)             Max: 22(0/3/CPU0)

        Priority: High
            Route Count:                 Added: 18          Deleted: 0          Modified:
 0
            FRR Coverage:                Routes: 9/9(100%)             Paths: 18/18(100%)

            IP Route Program Time:       Min: 11(0/1/CPU0)             Max: 12(0/2/CPU0)

            MPLS Label Program Time:     Min: 21(0/2/CPU0)             Max: 25(0/3/CPU0)

        Priority: Medium
            Route Count:                 Added: 18          Deleted: 0          Modified:
 0
            FRR Coverage:                Routes: 9/9(100%)             Paths: 18/18(100%)

            IP Route Program Time:       Min: 12(0/3/CPU0)             Max: 15(0/2/CPU0)

            MPLS Label Program Time:     Min: 22(0/2/CPU0)             Max: 26(0/3/CPU0)

        Priority: Low
            Route Count:                 Added: 21          Deleted: 0          Modified:
 0
            FRR Coverage:                Routes: 10/10(100%)           Paths: 21/21(100%)

            IP Route Program Time:       Min: 14(0/1/CPU0)             Max: 19(0/3/CPU0)

            MPLS Label Program Time:     Min: 28(0/1/CPU0)             Max: 33(0/2/CPU0)


Dijkstra Info:
```

```
Area: 0.0.0.0                   Run: 9
Trigger: Apr 10 23:07:01.562    Wait: 0             Start: 52           Duration: 0
```

storage-location

To specify where to store the extended routing-diagnostics that are collected when threshold exceeds, use the **storage-location** command in router-convergence configuration mode. To disable storing routing-diagnostics to a specific location, use the **no** form of this command.

storage-location [**diagnostics** *directory-path* | **diagnostics-size** *maximum-directory-size* | **reports** *directory-path* | **reports-size** *maximum-directory-size*]
no storage-location

Syntax Description

diagnostics	Specifies an absolute directory path for storing diagnostic reports.
directory-path	Specifies the path of the absolute directory for storing diagnostic reports.
diagnostics-size	Specifies the maximum size of diagnostics directory.
maximum-directory-size	Specified the size of the diagnostics directory. The range is 5% to 80%.
reports	Specifies an absolute directory path for storing reports.
directory-path	Specifies the path of the absolute directory for storing reports.
reports-size	Specifies the maximum size of the reports directory. The range is 5% to 80%.

Command Default

No default storage location. Mechanism is disabled.

Command Modes

Router-convergence configuration

Command History

Release	Modification
Release 4.2.0	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.

The storage location can be local disk or remote tftp space.

RCMD server can periodically archive reports in XML format for persistency. This mechanism is enabled when archival location is configured. Debug data collected in diagnostics mode is dumped to the configured diagnostic location (else it would get lost). When using local disk, the percentage of disk space to be used can be specified, and RCMD server will delete older reports on reaching the limit. Archival (specifically on local disk) is CPU intensive. Use a remote XML server to periodically collect reports from the router and archive on the server's local storage.

Task ID

Task ID	Operations
rcmd	read, write

Examples

This example shows how to configure storage location as *tftp://202.153.144.25/auto/tftp-chanvija-blr/rcmd/dump/reports* for reports and */harddisk:/rcmd_logs* for diagnostics:

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#router-convergence
RP/0/RSP0/CPU0:router(config-rcmd)#storage-location
RP/0/RSP0/CPU0:router(config-rcmd-store)#diagnostics /harddisk:/rcmd_logs
RP/0/RSP0/CPU0:router(config-rcmd-store)#reports
tftp://202.153.144.25/auto/tftp-chanvija-blr/rcmd/dump/reports
```

Related Commands

Command	Description
router-convergence, on page 1918	Configures route convergence monitoring and enters router convergence monitoring and diagnostics (rcmd) configuration mode.

track-external-routes

To enable tracking of external (Type-3/5/7) LSAs prefix monitoring, use the track-external-routes command in Router OSPF monitor-convergence configuration mode. To disable, tracking of external LSAs prefix monitoring, use the no form of this command.

track-external-routes
no track-external-routes

This command has no keywords or arguments.

Command Default

Route OSPF monitor-convergence

Command Modes

External LSAs prefix monitoring is disabled.

Command History

Release	Modification
Release 4.3.0	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
ospf	read, write

This example shows how to enable tracking of external LSAs prefix monitoring:

```
Router#configure
Router(config)#router ospf 100
Router(config-ospf)#monitor-convergence
Router(config-ospf-af-rcmd)#track-external-routes
```

track-summary-routes

To enable tracking of summary (inter-area) routes prefix monitoring, use the track-summary-routes command in Router OSPF monitor-convergence configuration mode. To disable tracking of summary router prefix monitoring, use the no form of this command.

track-summary-routes
no track-summary-routes

This command has no keywords or arguments.

Command Default

Router OSPF monitor-convergence

Command Modes

Summary routes prefix monitoring is disabled.

Command History

Release	Modification
Release 4.3.0	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
ospf	read, write

This example shows how to enable tracking of summary routes prefix monitoring:

```
Router#configure
Router(config)#router ospf 100
Router(config-ospf)#monitor-convergence
Router(config-ospf-af-rcmd)#track-summary-routes
```



Locator/ID Separation Protocol Commands

This module describes the commands used to configure and monitor Locator/ID Separation Protocol (LISP) on Cisco IOS XR.

- Locator/ID Separation Protocol on Cisco IOS XR, on page 1948
- address-family (LISP), on page 1949
- clear lisp vrf, on page 1951
- decapsulation filter rloc source, on page 1952
- eid-mtu, on page 1954
- eid-table, on page 1956
- etr, on page 1958
- etr accept-map-request-mapping, on page 1959
- etr map-cache-ttl, on page 1961
- etr map-server, on page 1963
- itr map-resolver, on page 1965
- locator reachability, on page 1967
- locator-set, on page 1968
- locator-table, on page 1970
- loc-reach-algorithm rloc-probing, on page 1972
- map-cache-limit, on page 1974
- map-cache, on page 1975
- map-request-source, on page 1977
- map-server rloc members distribute, on page 1978
- map-server rloc members modify-discovered {add | override} , on page 1980
- other-xtr-probe, on page 1982
- proxy-etr, on page 1984
- proxy-itr, on page 1986
- remote-rloc-probe, on page 1988
- router lisp, on page 1989
- show lisp decapsulation filter, on page 1991
- show lisp session, on page 1992
- show lisp site rloc members, on page 1993
- show lisp site, on page 1994
- solicit-map-request, on page 1996
- use-petr, on page 1998

Locator/ID Separation Protocol on Cisco IOS XR

Locator/ID Separation Protocol (LISP) is a simple, incremental, network-based protocol designed to implement separation of Internet addresses into Endpoint Identifiers (EIDs) and Routing Locators (RLOCs).

LISP stands for Locator/ID Separation Protocol and is a next-generation IP routing feature that creates a new paradigm in how IP addressing is assigned and interpreted by splitting the device identity, known as an endpoint identifier (EID), and its location, known as its routing locator (RLOC), into two different namespaces. Creating separate IP addresses for EID and RLOC functions yields several advantages, including improved scalability of the routing system through greater aggregation of RLOCs and improved multihoming efficiency and ingress traffic engineering. Hosts do not have to change IP addresses and therefore, no IP address numbering costs are involved with the LISP implementation.

LISP sites use IP addresses in the EID namespace to address hosts and in Domain Name System (DNS) in exactly the same way they are currently used. These addresses are not advertised within the non-LISP RLOC namespace (that is, the Internet), but instead are advertised by the LISP mapping services. The LISP site router supports the LISP functionality of Ingress Tunnel Router (ITR) and Egress Tunnel Router (ETR).

LISP is a pull model analogous to DNS and is massively scalable. LISP is address family agnostic and can be deployed incrementally.

LISP creates a Level of indirection with two namespaces: EID and RLOC. The EID (Endpoint Identifier) is the IP address of a host. The RLOC (Routing Locator) is the IP address of the LISP router for the host. EID-to-RLOC mapping is the distributed architecture that maps EIDs to RLOCs. The LISP Map Lookup is analogous to a DNS lookup. DNS resolves IP addresses for URLs. LISP resolves locators for queried identifiers or EID prefix.

LISP in Cisco IOS XR supports:

- Proxy Ingress Tunnel Router (PITR) and Proxy Egress Tunnel Router (PETR). PITR must be configured using map resolver (no ALT support).
- Default table support for EID and RLOC space.
- The **router lisp** command in global configuration mode enables LISP configuration mode.



Note The LISP command line interface, show commands output, and schema is to be changed in Cisco IOS XR Release 4.3.1 to be similar to the LISP command line interface on Cisco IOS.

For information on LISP Commands, see the *LISP Commands on Cisco IOS XR* module in the

address-family (LISP)

To enter Locator ID and separation protocol (LISP) address family configuration mode, use the **address-family** command in LISP configuration mode. To exit the LISP address family configuration mode, use the **no** form of this command.

address-family {**ipv4** | **ipv6**} **unicast**
no address-family {**ipv4** | **ipv6**} **unicast**

Syntax Description

ipv4	Selects IPv4 address family.
ipv6	Select IPv6 address family.
unicast	Selects unicast address prefixes.

Command Default

LISP address family configuration is disabled.

Command Modes

LISP configuration

Command History

Release	Modification
Release 4.3.0	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
lisp	read, write

This example shows how to enable IPv6 address family configuration for LISP:

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#router lisp
RP/0/RSP0/CPU0:router(config-lisp)#address-family ipv6 unicast
RP/0/RSP0/CPU0:router(config-lisp-afi)#
```

Related Commands

Command	Description
itr map-resolver, on page 1965	Configures an IPv4 or IPv6 locator address of the LISP Map-Resolver to be used by the ITR.
loc-reach-algorithm rloc-probing, on page 1972	Configures locator reachability algorithm, RLOC Probing, and determines reachability status for other iBGP peers through the IGP domain.

Command	Description
map-cache-limit, on page 1974	Configures the maximum limit of IPv4 LISP or IPv6 LISP map-cache entries allowed to be stored by the router.
map-cache, on page 1975	Configures a static IPv4 EID-to-RLOC or static IPv6 EID-to-RLOC mapping relationship and its associated traffic policy, or statically configures the packet handling behavior associated with a destination IPv4 EID-prefix or a destination IPv6 EID-prefix.
proxy-etr, on page 1984	Configures a router to act as an IPv4 or IPv6 LISP Proxy Egress Tunnel Router (PETR).
proxy-itr, on page 1986	Configures a router to act as an IPv4 or IPv6 LISP Proxy Ingress Tunnel Router (PITR).
router lisp, on page 1989	Enters Locator and ID Separation Protocol (LISP) configuration mode.

clear lisp vrf

To clear a Locator/ID Separation Protocol (LISP) reliable TCP transport session between an xTR and a Map-Server, use the **clear lisp vrf** command in the privileged EXEC mode.

clear lisp vrf *vrf-name* **session** {*peer-address* | *****}

Syntax Description

vrf-name	VRF instance. The transport session information for this VRF instance will be cleared. **Note** The vrf-name is a locator VRF, rather than an EID VRF. TCP sessions are formed per locator VRF. Hence, if you have shared mode configured where the locator VRF is default, then to clear the TCP sessions use the **clear lisp vrf default session {peer-address|*}** command.
session	Specifies that the reliable transport session for either the specified peer address or all transport sessions be cleared, based on your choice.
peer-address	IPv4 or IPv6 peer address. When you specify a peer-address, the TCP connection to the peer will be cleared.
*	Clears all LISP reliable transport sessions for this particular locator VRF. **Note** If you have multiple router LISP instances, other router LISP instances will not be affected.

Command Modes

EXEC mode

Command History

Release	Modification
Release 5.3.0	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

Task ID

Task ID	Operation
lisp	read, write

The following example shows how to clear all reliable TCP transport sessions with locator VRF v1 using the * option:

```
RP/0/RSP0/CPU0:router#clear lisp vrf v1 session *
```

decapsulation filter rloc source

To enable source Routing Locator (RLOC) address validation of Locator/ID Separation Protocol, (LISP)-encapsulated packets, use the **decapsulation filter rloc source** command in command in LISP configuration mode. This configures an xTR and a proxy-xTR to download decapsulation filter lists for source validation when decapsulating LISP packets.

To disable source RLOC address validation of LISP packets, use the **no** form of this command.

decapsulation filter rloc source {[**locator-set** *locator-set-name*] [**member**] }
no decapsulation filter rloc source

Syntax Description

locator-set *locator-set-name*	Specifies the full set of locators from which traffic will be accepted, both from other ITRs and PITRs.
member	Specifies that the registered RLOC membership list be automatically obtained from the Map-Server. **Note** Either or both of the **member** or **locator-set** keywords must be specified.

Command Default

Source RLOC address validation of LISP packets is disabled.

Command Modes

LISP configuration

Command History

Release	Modification
Release 5.3.0	This command was introduced.

Usage Guidelines

Configure this command on an xTR or a PxTR to enable LISP decapsulation filtering. When enabled, the source RLOC addresses of incoming LISP packets are validated against the 'member' filter list. RLOCs that match the filter list are decapsulated while those that do not are dropped.

When the member keyword is used, the registered RLOC membership list will be automatically obtained from the Map-Server. The **member** keyword enables the establishment of a reliable transport (TCP) session with configured map-servers, and downloads the decapsulation filter list maintained by the map-servers. When the **locator-set** *locator-set-name* keyword is used, the prefixes named in the locator-set are used, if included alone or added to the (downloaded) dynamic list when used in conjunction with the **member** keyword.

This option is used to add PITRs which do not register with a Map-Server and are thus not automatically included in the registered RLOC membership list.

- On an xTR, the TCP-based reliable transport session is established only after the UDP-based (normal) Map-Registration process successfully completes.
- On a PxTR, since this device does not (normally) register with a Map-Server, a “stub” (fake) Map-Registration configuration must be added to allow the establishment of the reliable transport session and the download of any filter lists. The Map-Server requires the PETR RLOC(s) to be included in a map-server rloc members modify-discovered add command to permit this session establishment.

- A (P)xTR normally communicates with multiple Map-Servers. However, in the event that all reliable transport session goes down, any existing (possibly stale) filter list will remain in use during a small window of time (several minutes), during which time the (P)xTR tries to re-establish the session(s) with the MS and refresh its membership.
- If no filter list can be downloaded, or the existing list times out, packets will be dropped. (fail closed.)
- If the xTR changes RLOCs (using DHCP for example), as soon as the RLOC is changed, the registration with the Map-Server is updated and the new registered RLOC is pushed to all “members” of this IID/VPN (event-driven).

Task ID

Task ID	Operation
lisp	read, write

This example shows how to configure decapsulation filter rloc source:

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#router lisp
RP/0/RSP0/CPU0:router(config-lisp)#decapsulation filter rloc source member
RP/0/RSP0/CPU0:router(config-lisp)#exit
```

eid-mtu

To configure MTU sizes for IPv4 or IPv6 LISP payloads, use the **eid-mtu** command in the LISP address family configuration mode. To remove the configured the eid-mtu size, use the **no** form of this command.

eid-mtu {**ipv4-locator** *ipv4_bytes* | **ipv6-locator** *ipv6_bytes*}
no eid-mtu

Syntax Description

ipv4-locator	Specifies the MTU for packets transported through IPv4 RLOC space.
ipv4_bytes	Specifies the MTU in bytes for packets transported through IPv4 RLOC space. The value ranges from 68 to 65535.
ipv6-locator	Specifies the MTU for packets transported through IPv6 RLOC space.
ipv6_bytes	Specifies the MTU in bytes for packets transported through IPv6 RLOC space. The value ranges from 1280 to 65535.

Command Default

None

Command Modes

LISP IPv4 address family

LISP IPv6 address family

Command History

Release	Modification
Release 4.3.1	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	Operation
lisp	read, write

This example shows how to configure MTU size of 1300 bytes for the IPv6 LISP payloads:

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#router lisp
```

```
RP/0/RSP0/CPU0:router(config-lisp)#address-family ipv6 unicast
RP/0/RSP0/CPU0:router(config-lisp-af)#eid-mtu ipv4-locator 1320 ipv6-locator 1300
```

Related Commands

Command	Description
router lisp, on page 1989	Enters Locator and ID Separation Protocol (LISP) configuration mode.
address-family (LISP), on page 1949	Enters Locator ID and separation protocol (LISP) address family configuration mode.

eid-table

To configure a Locator ID Separation Protocol (LISP) instance ID for association with a virtual routing and forwarding (VRF) table or default table through which the endpoint identifier (EID) address space is reachable, use the **eid-table** command in the LISP configuration mode. To remove this association, use the **no** form of this command.

eid-table {**default** | [**vrf** **vrf_name**]} **instance-id** *instance_id*
no **eid-table** {**default** | **vrf** **vrf_name**} **instance-id** *instance_id*

Syntax Description

default	Selects the default (global) routing table for association with the configured instance ID.
vrf	Selects the specified VRF table for association with the configured instance ID.
vrf_name	Specifies the name of the VRF.
instance	Specifies the instance ID to be associated with this EID table.
instance_id	Specifies the instance ID value. This value ranges between 0 and 16777215.

Command Default

A router configured for LISP associates the default table with instance ID 0.

Command Modes

LISP configuration

Command History

Release	Modification
Release 4.3.1	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.

The **eid-table** command is used to associate a LISP instance ID with either the default routing table, or a VRF table through which its EID address space is reachable. When a LISP instance ID is specified, LISP Map Registration (control plane) messages include this instance ID along with the associated EID prefixes upon registering and LISP data plane packets include this instance ID in the LISP header.

Task ID

Task ID	Operation
lisp	read, write

This example shows how to configure LISP instance ID for association with a virtual routing and forwarding (VRF) table:

```
RP/0/RSP0/CPU0:router#configure
```

```
RP/0/RSP0/CPU0:router(config)#router lisp
RP/0/RSP0/CPU0:router(config-lisp)#eid-table vrf vrf1 instance-id 45
```

Related Commands

Command	Description
locator-table, on page 1970	Specifies the RLOC table.
router lisp, on page 1989	Enters Locator and ID Separation Protocol (LISP) configuration mode.
itr map-resolver, on page 1965	Configures an IPv4 or IPv6 locator address of the LISP Map-Resolver to be used by the ITR.
map-cache-limit, on page 1974	Configures the maximum limit of IPv4 LISP or IPv6 LISP map-cache entries allowed to be stored by the router.
map-cache, on page 1975	Configures a static IPv4 EID-to-RLOC or static IPv6 EID-to-RLOC mapping relationship and its associated traffic policy, or statically configures the packet handling behavior associated with a destination IPv4 EID-prefix or a destination IPv6 EID-prefix.

etr

To enable Egress Tunnel Router functionality, use the **etr** command in the LISP address family configuration mode. To disable the ETR functionality, use the **no** form of this command.

etr
no etr

Syntax Description

This command has no keywords or arguments.

Command Default

ETR functionality is not enabled by default.

Command Modes

LISP IPv4 address family

LISP IPv6 address family

Command History

Release	Modification
Release 4.3.1	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	Operation
lisp	read, write

The example shows how to turn on ETR functionality in all eid-tables, unless it is explicitly disabled.

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#router lisp
RP/0/RSP0/CPU0:router(config-lisp)#address-family ipv6 unicast
RP/0/RSP0/CPU0:router(config-lisp-af)#etr
```

Related Commands

Command	Description
router lisp, on page 1989	Enters Locator and ID Separation Protocol (LISP) configuration mode.
address-family (LISP), on page 1949	Enters Locator ID and separation protocol (LISP) address family configuration mode.

etr accept-map-request-mapping

To configure an Egress Tunnel Router (ETR) to cache to cache EID-to-RLOC mapping data contained, that ITRs may attach to a map-request message, use the **etr accept-map-request-mapping** command in the LISP address family configuration mode. To remove this functionality, use the **no** form of this command.

etr accept-map-request-mapping [**verify**]
no etr accept-map-request-mapping

Syntax Description

verify	(Optional) Specifies that mapping data should be cached but not used for forwarding packets until the ETR can send its own map request to one of the locators from the mapping data record and receive a map reply with the same data in response.

Command Default

No caching of mapping data in a map-request message.

Command Modes

LISP IPv4 address family

LISP IPv6 address family

Command History

Release	Modification
Release 4.3.1	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.

If an ETR receives a map-request message that contains mapping data for the invoking IPv6 source-EID's packet, then the ETR, by default, ignores the mapping data. However, if you configure the **etr accept-map-request-mapping** command, the ETR caches the mapping data in its map cache and immediately uses it for forwarding packets.

Task ID

Task ID	Operation
lisp	read, write

This example shows how to configure **etr accept-map-request-mapping** command in the LISP IPv6 address family configuration mode:

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#router lisp
```

```
RP/0/RSP0/CPU0:router(config-lisp)#address-family ipv6 unicast
RP/0/RSP0/CPU0:router(config-lisp-af)#etr accept-map-request-mapping verify
```

Related Commands

Command	Description
router lisp, on page 1989	Enters Locator and ID Separation Protocol (LISP) configuration mode.
address-family (LISP), on page 1949	Enters Locator ID and separation protocol (LISP) address family configuration mode.
etr, on page 1958	Enables Egress Tunnel Router (ETR) functionality.

etr map-cache-ttl

To configure the time-to-live (TTL) value inserted into LISP IPv4 or IPv6 map-reply messages, use the **etr map-cache-ttl** command in the LISP address family configuration mode. To remove the configured TTL value and return to the default value, use the **no** form of this command.

etr map-cache-ttl *minutes*
no etr map-cache-ttl *minutes*

Syntax Description

minutes	Specifies the value, in minutes, to be inserted in the TTL field in map-reply messages. The value ranges from 60 to 10080 minutes.

Command Default

The default TTL value is 1440 minutes (24 hours).

Command Modes

LISP IPv4 address family

LISP IPv6 address family

Command History

Release	Modification
Release 4.3.1	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.

Use this command to change the default value associated with the Time-to-Live (TTL) field in map-reply messages. Entering this command changes the default TTL that remote ITRs will cache and use for your site's endpoint identifier (EID) prefix. The default value is 1440 minutes (24 hours), and the minimum value is 60 minutes.

Task ID

Task ID	Operation
lisp	read, write

This example shows how to configure **etr map-cache-ttl** command with TTL value of 60 minutes:

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#router lisp
RP/0/RSP0/CPU0:router(config-lisp)#address-family ipv6 unicast
RP/0/RSP0/CPU0:router(config-lisp-af)#etr map-cache-ttl 60
```

Related Commands

Command	Description
router lisp, on page 1989	Enters Locator and ID Separation Protocol (LISP) configuration mode.

Command	Description
address-family (LISP), on page 1949	Enters Locator ID and separation protocol (LISP) address family configuration mode.
etr, on page 1958	Enables Egress Tunnel Router (ETR) functionality.

etr map-server

To configure the options related to the etr map-server (MS) such as locator, authentication key and whether or not the map server is allowed to respond on the ETR's behalf to map-requests (proxy-reply option), use the **etr map-server** command in the LISP address family configuration mode. To remove the configured options, use the **no** form of this command.

etr map-server *IP* {**key** [**clear** | **encrypted**] **LINE** | **proxy-reply**}
no etr map-server

Syntax Description

ip_address	Specifies the IPv4 or IPv6 address of the map server.
key	Specifies that an authentication key will follow either encrypted or unencrypted option.
LINE	Specifies the key (either encrypted or unencrypted).
clear	Specifies the UNENCRYPTED (cleartext) key.
encrypted	Specifies that an ENCRYPTED key will follow.
proxy-reply	Specifies that the map-server with the specified RLOC is allowed to respond to map-requests on behalf of the ETR.

Command Default

LISP map server locator addresses are not configured by default.

Command Modes

LISP IPv4 address family

LISP IPv6 address family

Command History

Release	Modification
Release 4.3.1	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	Operation
lisp	read, write

This example configures 2001:0db8::1 as map-server, specifying that *supers3cr3tpassw0rd* will be used as the shared secret for HMAC calculations, and that this map-server may respond to map-requests on behalf of the ETR.

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#router lisp
```

```
RP/0/RSP0/CPU0:router(config-lisp)#address-family ipv6 unicast
RP/0/RSP0/CPU0:router(config-lisp-af)#etr map-server 2001:0db8::1 key clear
supers3cr3tpassw0rd
RP/0/RSP0/CPU0:router(config-lisp-af)#etr map-server 2001:0db8::1 proxy-reply
```

Related Commands

Command	Description
router lisp, on page 1989	Enters Locator and ID Separation Protocol (LISP) configuration mode.
address-family (LISP), on page 1949	Enters Locator ID and separation protocol (LISP) address family configuration mode.
etr, on page 1958	Enables Egress Tunnel Router (ETR) functionality.

itr map-resolver

To configure an IPv4 or IPv6 locator address of the LISP Map-Resolver to be used by the ITR, when sending Map-Requests for IPv4 EID-to-RLOC mapping resolution, use the **itr map-resolver** command in LISP address family configuration mode. To remove the configured locator address of the LISP Map-Resolver, use the **no** form of this command.

itr map-resolver *map-resolver-address*
no itr map-resolver *map-resolver-address*

Syntax Description

map-resolver-address	Specifies the IPv4 or IPv6 locator addresses of the Map-Resolver.

Command Default

No LISP Map-Resolver locator address is configured.

Command Modes

LISP IPv4 address family configuration

LISP IPv6 address family configuration

Command History

Release	Modification
Release 4.3.0	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
lisp	read, write

This example configures an ITR to use the Map-Resolver located at 10.2.3.4 when sending its Map-Request messages:

```
RP/0/RSP0/CPU0:router(config)#router lisp
RP/0/RSP0/CPU0:router(config-lisp)#address-family ipv6 unicast
RP/0/RSP0/CPU0:router(config-lisp-afi)#itr map-resolver 10.2.3.4
```

This example configures and ITR to use the Map-Resolver located at 2001:DB8:0A::1 when sending its Map-Request messages:

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#router lisp
RP/0/RSP0/CPU0:router(config-lisp)#address-family ipv6 unicast
```

```
RP/0/RSP0/CPU0:router(config-lisp-afi)#itr map-resolver 2001:DB8:0A::1
```

Related Commands

Command	Description
address-family (LISP), on page 1949	Enters Locator ID and separation protocol (LISP) address family configuration mode.
loc-reach-algorithm rloc-probing, on page 1972	Configures locator reachability algorithm, RLOC Probing, and determines reachability status for other iBGP peers through the IGP domain.
map-cache-limit, on page 1974	Configures the maximum limit of IPv4 LISP or IPv6 LISP map-cache entries allowed to be stored by the router.
map-cache, on page 1975	Configures a static IPv4 EID-to-RLOC or static IPv6 EID-to-RLOC mapping relationship and its associated traffic policy, or statically configures the packet handling behavior associated with a destination IPv4 EID-prefix or a destination IPv6 EID-prefix.
proxy-etr, on page 1984	Configures a router to act as an IPv4 or IPv6 LISP Proxy Egress Tunnel Router (PETR).
proxy-itr, on page 1986	Configures a router to act as an IPv4 or IPv6 LISP Proxy Ingress Tunnel Router (PITR).
router lisp, on page 1989	Enters Locator and ID Separation Protocol (LISP) configuration mode.

locator reachability

To configure the conditions that must be met for a routing locator (RLOC) to be treated as being reachable, use the **locator reachability** command in the LISP address family configuration mode. To remove the configuration, use the **no** form of the command.

locator reachability exclude-default
no locator reachability exclude-default

Syntax Description

exclude-default	If a remote RLOC is reached via the default route, then it is treated as unreachable.

Command Default

Command Modes

LISP IPv4 address family

LISP IPv6 address family

Command History

Release	Modification
Release 4.3.1	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	Operation
lisp	read, write

This example shows how to configure the **locator reachability** command:

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#router lisp
RP/0/RSP0/CPU0:router(config-lisp)#address-family ipv6 unicast
RP/0/RSP0/CPU0:router(config-lisp-af)#locator reachability exclude-default
```

Related Commands

Command	Description
loc-reach-algorithm rloc-probing, on page 1972	Configures locator reachability algorithm, RLOC Probing, and determines reachability status for other iBGP peers through the IGP domain.
address-family (LISP), on page 1949	Enters Locator ID and separation protocol (LISP) address family configuration mode.
router lisp, on page 1989	Enters Locator and ID Separation Protocol (LISP) configuration mode.

locator-set

To configure a named locator set, use the **locator-set** command in the in LISP configuration mode. To disable a the locator-set, use the **no** form of the command.

locator-set *name* {**ip_address** {**priority** *priority_value* | **weight** *weight_value*} | **auto-discover-rlocs**}
no locator-set *name*

Syntax Description

name	Specifies the name of the locator set.
ip_address	Specifies the RLOC IP address of Loopback or other Egress Tunnel Router (ETR) interfaces.
priority	Configures the preferred locators from the locator set. When multiple locators have the same priority traffic may be load-balanced across them.
priority_value	Specifies the value of the priority assigned to the RLOC. The value ranges from 0 to 255. A lower value indicates a higher priority.
weight	Specifies how to determine to load-share traffic between multiple locators when the priorities assigned to multiple locators are the same.
weight_value	Specifies the value of the percentage of traffic to be load-shared. The value ranges from 0 to 100.
auto-discover-rlocs	Configures the Egress Tunnel Router (ETR) to discover the locators of all routers configured to function as both an ETR and an Ingress Tunnel Router (ITR)--such routers are referred to as xTRs--in ETRs LISP site when the site uses multiple xTRs and each xTR is configured to use DHCP-learned locators or configured with only its own locators.

Command Default

Command Modes

LISP configuration

Command History

Release	Modification
Release 4.3.1	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.

A locator set is a named group of Routing LOCators (RLOCs). It can be used in conjunction with the **database-mapping** and **map-cache** configuration commands.

Task ID

Task ID	Operation
lisp	read, write

This example shows how to configure the locator-set priority such that 2001:1::2 would have the highest priority and receive all the traffic. 2001:2::3 and 64.10.10.5 would each receive 50% of the traffic if 2001:1::2 is unreachable.

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#router lisp
RP/0/RSP0/CPU0:router(config-lisp)#locator-set loc1
RP/0/RSP0/CPU0:router(config-lisp)#2001:1::2 priority 1 weight 100
RP/0/RSP0/CPU0:router(config-lisp)#2001:2::3 priority 2 weight 50
RP/0/RSP0/CPU0:router(config-lisp)#64.10.10.5 priority 2 weight 50
```

Related Commands

Command	Description
address-family (LISP), on page 1949	Enters Locator ID and separation protocol (LISP) address family configuration mode.
loc-reach-algorithm rloc-probing, on page 1972	Configures locator reachability algorithm, RLOC Probing, and determines reachability status for other iBGP peers through the IGP domain.
router lisp, on page 1989	Enters Locator and ID Separation Protocol (LISP) configuration mode.

locator-table

To associate a virtual routing and forwarding (VRF) table through which the routing locator address space is reachable to a router Locator ID Separation Protocol (LISP) instantiation, use the **locator-table** command in LISP configuration mode. To remove this association, use the **no** form of this command.

locator-table *name* [**default** | **vrf** *vrf_name*]
no locator-table *name*

Syntax Description

default	Selects the default (global) routing table for association with the routing locator address space.
vrf	Selects the routing table for the specified VRF name for association with the routing locator address space.
vrf_name	Specifies the name of the VRF.

Command Default

None

Command Modes

LISP configuration

Command History

Release	Modification
Release 4.3.1	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.

Note The locator-table configuration is mandatory for LISP to function.

When a LISP device is deployed in a multitenant (virtualized) network environment with segmented routing locator (RLOC) address space, separate router LISP instantiations are required for each locator address space. Separate instantiations are created by including the optional id entry with the **router lisp** command. Each router LISP instantiation is considered to be standalone and must be associated with an RLOC address space. The **locator-table** command is used to associate a VRF table through which the routing locator address space is reachable to a router LISP instantiation. All necessary LISP components used in the operation of that particular router LISP instantiation, (for example, map server, map resolver, proxy ingress tunnel router (PITR), proxy egress tunnel router (PETR), and other routers that function as both egress and ingress tunnel routers, also known as xTRs) must be reachable via the routing locator address space referred to by the **locator-table** command.

Task ID

Task ID	Operation
lisp	read, write

This example shows how to associate a VRF table to a LISP instantiation:

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#router lisp
RP/0/RSP0/CPU0:router(config-lisp)#locator-table  mplsvpn
```

Related Commands

Command	Description
address-family (LISP), on page 1949	Enters Locator ID and separation protocol (LISP) address family configuration mode.
loc-reach-algorithm rloc-probing, on page 1972	Configures locator reachability algorithm, RLOC Probing, and determines reachability status for other iBGP peers through the IGP domain.
router lisp, on page 1989	Enters Locator and ID Separation Protocol (LISP) configuration mode.

loc-reach-algorithm rloc-probing

To configure locator reachability algorithm, RLOC Probing, and to determine reachability status for other iBGP peers through the IGP domain, use the **loc-reach-algorithm rloc-probing** command in LISP configuration mode. To disable the locator reachability algorithm, use the **no** form of this command .

loc-reach-algorithm rloc-probing
no loc-reach-algorithm rloc-probing

Syntax Description

This command has no keywords or arguments.

Command Default

The locator reachability algorithm rloc-probing is disabled.

Command Modes

LISP configuration

Command History

Release	Modification
Release 4.3.0	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
lisp	read, write

This example shows how to configure the location reachability algorithm RLOC probing:

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#router lisp
RP/0/RSP0/CPU0:router(config-lisp)#loc-reach-algorithm rloc-probing
```

Related Commands

Command	Description
address-family (LISP), on page 1949	Enters Locator ID and separation protocol (LISP) address family configuration mode.
itr map-resolver, on page 1965	Configures an IPv4 or IPv6 locator address of the LISP Map-Resolver to be used by the ITR.
map-cache-limit, on page 1974	Configures the maximum limit of IPv4 LISP or IPv6 LISP map-cache entries allowed to be stored by the router.
map-cache, on page 1975	Configures a static IPv4 EID-to-RLOC or static IPv6 EID-to-RLOC mapping relationship and its associated traffic policy, or statically configures the

Command	Description
	packet handling behavior associated with a destination IPv4 EID-prefix or a destination IPv6 EID-prefix.
proxy-etr, on page 1984	Configures a router to act as an IPv4 or IPv6 LISP Proxy Egress Tunnel Router (PETR).
proxy-itr, on page 1986	Configures a router to act as an IPv4 or IPv6 LISP Proxy Ingress Tunnel Router (PITR).
router lisp, on page 1989	Enters Locator and ID Separation Protocol (LISP) configuration mode.

map-cache-limit

To configure the maximum limit of IPv4 LISP or IPv6 LISP map-cache entries allowed to be stored by the router, use the **map-cache-limit** command in LISP address family configuration mode. To remove the configured map-cache limit, use the **no** form of this command.

map-cache-limit *map-cache-size*
no map-cache-limit *map-cache-size*

Syntax Description

map-cache-size	Specifies the map cache size value. Range is 1 to 65535.

Command Default

Map cache size: 1000

Command Modes

LISP IPv4 address family

LISP IPv6 address family

Command History

Release	Modification
Release 4.3.0	This command was introduced.

Usage Guidelines

None

Task ID

Task ID	Operation
lisp	read, write

This example configures a lisp cache-limit of 2000 entries:

```
Router(config-lisp-afi)#map-cache-limit 2000
```

map-cache

To configure a static IPv4 EID-to-RLOC or static IPv6 EID-to-RLOC mapping relationship and its associated traffic policy, or to statically configure the packet handling behavior associated with a destination IPv4 EID-prefix or a destination IPv6 EID-prefix, use the **map-cache** command in LISP address family configuration mode. To remove the configuration, use the **no** form of this command.

map-cache *destination-EID-prefix / prefix-length* {**action** {**drop** | **map-request** | **native-forward**} | **locator** *locator-address* **priority** *priority-value* **weight** *weight-value*}
no map-cache *destination-EID-prefix / prefix-length* {**action** {**drop** | **map-request** | **native-forward**} | **locator** *locator-address* **priority** *priority-value* **weight** *weight-value*}

Syntax Description

destination-EID-prefix / prefix-length	Specifies the destination IPv4/IPv6 EID-prefix/prefix-length.
action	Specifies the non LISP forwarding action for the EID prefix.
drop	Selects drop action for the EID-prefix.
map-request	Generates a map-request for the EID-prefix.
native-forward	Specifies to natively forward EID-prefix.
locator *locator-address*	The IPv4 or IPv6 Routing Locator (RLOC) associated with the EID-prefix/prefix-length.
priority *priority-value*	The priority (value between 0 and 255) assigned to the RLOC. When multiple locators have the same priority they may be used in load-shared fashion. A lower value indicates a higher priority.
weight *weight-value*	The weight (value between 0 and 100) assigned to the locator. Used in order to determine how to load-share traffic between multiple locators when the priorities assigned to multiple locators are the same. The value represents the percentage of traffic to be load-shared.

Command Default

No IPv6 EID-to-RLOC mapping relationships or static IPv6 EID-to-RLOC mapping destinations are configured.

Command Modes

LISP IPv4 address family

LISP IPv6 address family

Command History

Release	Modification
Release 4.3.0	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
lisp	read, write

configures a destination EID-to-RLOC mapping and associated traffic policy for the IPv6 EID-prefix block 2001:DB8:BB::/48. The locator for this EID-prefix block is 2001:DB8:0A::1 and the traffic policy for this locator has a priority of 1 and a weight of 100.

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#router lisp
RP/0/RSP0/CPU0:router(config-lisp)#address-family ipv6 unicast
RP/0/RSP0/CPU0:router(config-lisp-afi)#map-cache
RP/0/RSP0/CPU0:router(config-lisp-afi-map-cache)#2001:db8:bb::/48 locator 2001:db8:a::1
 priority 1 weight 100
```

Related Commands

Command	Description
address-family (LISP), on page 1949	Enters Locator ID and separation protocol (LISP) address family configuration mode.
itr map-resolver, on page 1965	Configures an IPv4 or IPv6 locator address of the LISP Map-Resolver to be used by the ITR.
loc-reach-algorithm rloc-probing, on page 1972	Configures locator reachability algorithm, RLOC Probing, and determines reachability status for other iBGP peers through the IGP domain.
map-cache-limit, on page 1974	Configures the maximum limit of IPv4 LISP or IPv6 LISP map-cache entries allowed to be stored by the router.
proxy-etr, on page 1984	Configures a router to act as an IPv4 or IPv6 LISP Proxy Egress Tunnel Router (PETR).
proxy-itr, on page 1986	Configures a router to act as an IPv4 or IPv6 LISP Proxy Ingress Tunnel Router (PITR).
router lisp, on page 1989	Enters Locator and ID Separation Protocol (LISP) configuration mode.

map-request-source

To configure inner header source address to be used as the source address for Locator/ID Separation Protocol (LISP) map-request messages, use the **map-request-source** command in LISP address family configuration mode. To remove the configured map-request source address, use the **no** form of this command.

map-request-source *ip_address*
no map-request-source

Syntax Description

ip_address	Specifies the IPv4 or IPv6 source address of the inner header in the map-request message.

Command Default

Command Modes

LISP IPv4 address family

LISP IPv6 address family

Command History

Release	Modification
Release 4.3.1	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	Operation
lisp	read, write

This example shows how to configure the IPv6 source address to be used as the source address for LISP map-request messages:

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#router lisp
RP/0/RSP0/CPU0:router(config-lisp)#address-family ipv6 unicast
RP/0/RSP0/CPU0:router(config-lisp-af)#map-request-source 4:5::6
```

Related Commands

Command	Description
itr map-resolver, on page 1965	Configures an IPv4 or IPv6 locator address of the LISP Map-Resolver to be used by the ITR.
address-family (LISP), on page 1949	Enters Locator ID and separation protocol (LISP) address family configuration mode.
router lisp, on page 1989	Enters Locator and ID Separation Protocol (LISP) configuration mode.

map-server rloc members distribute

To enable Map-Servers to distribute a membership list of Routing Locators (RLOCs) to participating xTRs, use the **map-server rloc members distribute** command in LISP configuration mode. To disable Map-Servers from distributing a membership list of RLOCs to participating xTRs, use the no form of the command.

map-server rloc members distribute
no map-server rloc members distribute

Syntax Description

This command has no keywords or arguments.

Command Default

Map-Servers are not enabled to distribute a membership list of RLOCs to xTRs.

Command Modes

LISP configuration

Command History

Release	Modification
Release 5.3.0	This command was introduced.

Usage Guidelines

When an ETR or PETR decapsulates LISP-encapsulated packets, it normally does so without consideration of the outer (RLOC) header source address. In networking environments where the source address can be trusted, it may be useful for devices to consider the source address of the LISP packet prior to decapsulation. A Map-Server can be configured to dynamically create, maintain, and distribute decapsulation filter lists, on a per instance-ID basis, to appropriate LISP devices using the map-server rloc members distribute command in site configuration mode. When configured:

- The Map-Server allows the establishment of TCP-based LISP reliable transport sessions with appropriate xTRs
- The Map-Server creates/maintains lists (per-IID) of LISP site RLOCs based on RLOC addresses of registered LISP sites
- The Map-Server pushes/updates filters lists over the reliable transport mechanism to established devices



Note Data-plane security is enabled by the use of the **map-server roc members distribute** command. The optional command **map-server rloc members modified-discovered [add | override]** is used to append to or override the dynamically maintained RLOC filter list.

This feature is used in conjunction with the decapsulation filter rloc source command, configured on (P)xTR devices which are performing the decapsulation.

Task ID

Task ID	Operation
lisp	read, write

The following example shows how to enable Map-Servers to distribute an updated list of EID prefixes to xTRs:

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#router lisp
RP/0/RSP0/CPU0:router(config-lisp)#map-server rloc members distribute
RP/0/RSP0/CPU0:router(config-lisp)#exit
```

map-server rloc members modify-discovered {add | override}

To enable a Map-Server to add to, or replace, the list of discovered Routing Locator (RLOC) addresses through a specified locator set, use the **map-server rloc members modify-discovered** command in LISP EID-table configuration mode. To disable the option to modify the list of discovered RLOC addresses, use the no form of the command.

map-server rloc members modify-discovered {**add** | **override**} **locator-set** *locator-set-name*
no map-server rloc members modify-discovered

Syntax Description

add	Adds RLOC addresses in the specified locator set to the list of discovered RLOC addresses.
override	Replaces automatically discovered list of RLOC addresses with the RLOC addresses in the specified locator-set. When you use this option, the list of RLOC addresses discovered by the Map-Server is completely removed.
locator-set *locator-set-name*	Specifies a locator set. The locator set contains RLOC addresses that are previously configured.

Command Default

The option to modify the automatically discovered list of RLOC addresses is disabled.

Command Modes

LISP EID-table configuration

Command History

Release	Modification
Release 5.3.0	This command was introduced.

Usage Guidelines

When an ETR or PETR is configured to source-filter LISP-encapsulated packets prior to decapsulation, it may be useful to add to, or in some cases, override this filter list, and this can be accomplished by configuring the map-server rloc members modify-discovered command in EID-table configuration mode. Uses may include:

- When a PxTR is included in the architecture, the PITR LISP-encapsulates packets to an ETR – and the ETR must therefore include the PITR RLOC in its decapsulation filter list. Since PITRs do not register with Map-Servers, their RLOCs are not automatically included in the decapsulation filter list and must be added via configuration using this command.
- A PETR can also be configured to filter upon decapsulation, but again, because a PETR does not register with a Map-Server, it needs a way to obtain the decapsulation filter list. The add form of this command includes the mechanisms to establish the reliable transport session with the Map-Server for obtaining the decapsulation filter list on the PETR.
- For diagnostic/troubleshooting reasons, it may be useful to (temporarily) override the entire decapsulation filter list.

Task ID

Task ID	Operation
lisp	read, write

The following example shows how to enable Map-Servers to distribute an updated list of EID prefixes to xTRs:

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#router lisp
RP/0/RSP0/CPU0:router(config-lisp)#eid-table vrf cust-A instance-id 1
RP/0/RSP0/CPU0:router(config-lisp-eid-table)#map-server rloc members modify-discovered add
 locator-set PTR_set
RP/0/RSP0/CPU0:router(config-lisp-eid-table)#exit
```

other-xtr-probe

To configure parameters for probing of other xTRs *site-local* routing locators (RLOCs), use the **other-xtr-probe** command in the LISP configuration mode. To return to the default setting, use the **no** form of this command.

other-xtr-probe period *seconds*
no other-xtr-probe period *seconds*

Syntax Description

period	Configures the site-local RLOC probing period.
seconds	Specifies the value of the probing period in seconds. The range is between 5 to 900.

Command Default

Probing of site-local RLOCs is enabled by default and cannot be disabled. The default interval is 30 seconds.

Command Modes

LISP configuration

Command History

Release	Modification
Release 4.3.1	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.

Use the **other-xtr-probe** command to change the probe interval for sending RLOC probes to all site-other RLOCs.

This command probes site-local rlocs, whereas rloc-probing probes remote RLOCs. The latter can be turned on or off with loc-reach-algorithm rloc-probing. Remote-rloc-probe also probes remote RLOCs.

Note This functionality is enabled by default and cannot be disabled. The default interval is 30 seconds. Use the **show run | include other-xtr-probe** command to display the configured interval. When an output value is displayed, the value is configured for something other than the default value. When no output is displayed, it is configured for the default.

Task ID

Task ID	Operation
lisp	read, write

This example shows how to configure the probing period of 86 seconds for probing RLOCs:

```
RP/0/RSP0/CPU0:router#configure
```

```
RP/0/RSP0/CPU0:router(config)#router lisp
RP/0/RSP0/CPU0:router(config-lisp)#other-xtr-probe period 86
```

Related Commands

Command	Description
remote-rloc-probe, on page 1988	Configures parameters for probing of remote local routing locators (RLOCs)
address-family (LISP), on page 1949	Enters Locator ID and separation protocol (LISP) address family configuration mode.
router lisp, on page 1989	Enters Locator and ID Separation Protocol (LISP) configuration mode.

proxy-etr

To configure a router to act as an IPv4 or IPv6 LISP Proxy Egress Tunnel Router (PETR), use the **proxy-etr** command in LISP address family configuration mode. To remove LISP PETR functionality, use the **no** form of this command.

proxy-etr
no proxy-etr

This command has no keywords or arguments.

Command Default

ETR functionality is disabled.

Command Modes

LISP IPv4 address family

LISP IPv6 address family

Command History

Release	Modification
Release 4.3.0	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
lisp	read, write

This example shows how to configure PETR functionality under LISP IPv6 on the router:

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#router lisp
RP/0/RSP0/CPU0:router(config-lisp)#address-family ipv6 unicast
RP/0/RSP0/CPU0:router(config-lisp-afi)#proxy-etr
```

Related Commands

Command	Description
address-family (LISP), on page 1949	Enters Locator ID and separation protocol (LISP) address family configuration mode.
itr map-resolver, on page 1965	Configures an IPv4 or IPv6 locator address of the LISP Map-Resolver to be used by the ITR.
loc-reach-algorithm rloc-probing, on page 1972	Configures locator reachability algorithm, RLOC Probing, and determines reachability status for other iBGP peers through the IGP domain.

Command	Description
map-cache-limit, on page 1974	Configures the maximum limit of IPv4 LISP or IPv6 LISP map-cache entries allowed to be stored by the router.
map-cache, on page 1975	Configures a static IPv4 EID-to-RLOC or static IPv6 EID-to-RLOC mapping relationship and its associated traffic policy, or statically configures the packet handling behavior associated with a destination IPv4 EID-prefix or a destination IPv6 EID-prefix.
proxy-itr, on page 1986	Configures a router to act as an IPv4 or IPv6 LISP Proxy Ingress Tunnel Router (PITR).
router lisp, on page 1989	Enters Locator and ID Separation Protocol (LISP) configuration mode.

proxy-itr

To configure a router to act as an IPv4 or IPv6 LISP Proxy Ingress Tunnel Router (PITR), use the **proxy-itr** command in LISP address family configuration mode. To remove LISP PITR functionality, use the **no** form of this command.

proxy-itr *IPv4-source-locator-address*
noproxy-itr *IPv4-source-locator-address*

Syntax Description

IPv4-source-locator-address	Specifies the IPv4 source locator for PITR.

Command Default

PITR functionality is disabled.

Command Modes

LISP IPv4 address family

LISP IPv6 address family

Command History

Release	Modification
Release 4.3.0	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
lisp	read, write

This example shows how to configure LISP PITR functionality on the router:

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#router lisp
RP/0/RSP0/CPU0:router(config-lisp)#address-family ipv6 unicast
RP/0/RSP0/CPU0:router(config-lisp-afi)#proxy-itr 10.2.3.4
```

Related Commands

Command	Description
address-family (LISP), on page 1949	Enters Locator ID and separation protocol (LISP) address family configuration mode.
itr map-resolver, on page 1965	Configures an IPv4 or IPv6 locator address of the LISP Map-Resolver to be used by the ITR.

Command	Description
loc-reach-algorithm rloc-probing, on page 1972	Configures locator reachability algorithm, RLOC Probing, and determines reachability status for other iBGP peers through the IGP domain.
map-cache-limit, on page 1974	Configures the maximum limit of IPv4 LISP or IPv6 LISP map-cache entries allowed to be stored by the router.
map-cache, on page 1975	Configures a static IPv4 EID-to-RLOC or static IPv6 EID-to-RLOC mapping relationship and its associated traffic policy, or statically configures the packet handling behavior associated with a destination IPv4 EID-prefix or a destination IPv6 EID-prefix.
proxy-etr, on page 1984	Configures a router to act as an IPv4 or IPv6 LISP Proxy Egress Tunnel Router (PETR).
router lisp, on page 1989	Enters Locator and ID Separation Protocol (LISP) configuration mode.

remote-rloc-probe

To configure parameters for probing of remote local routing locators (RLOCs), use the **remote-rloc-probe** command in the LISP configuration mode. To return to the default setting, use the **no** form of this command.

remote-rloc-probe on-route-change
no remote-rloc-probe on-route-change

Syntax Description

on-route-change	Specifies the probing of the trigger on routing changes for remote RLOCs.

Command Default

None

Command Modes

LISP configuration

Command History

Release	Modification
Release 4.3.1	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	Operation
lisp	read, write

This example shows how to configure the **remote-rloc-probe** command for remote RLOCs:

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#router lisp
RP/0/RSP0/CPU0:router(config-lisp)#remote-rloc-probe on-route-change
```

Related Commands

Command	Description
address-family (LISP), on page 1949	Enters Locator ID and separation protocol (LISP) address family configuration mode.
router lisp, on page 1989	Enters Locator and ID Separation Protocol (LISP) configuration mode.

router lisp

To enter Locator and ID Separation Protocol (LISP) configuration mode, use the **router lisp** command in global configuration mode. To remove all LISP configurations and exit the LISP configuration mode, use the **no** form of this command.

router lisp *process-number*
no router lisp

Syntax Description

process-number	Specifies the router LISP process number. The range is 0-15.

Syntax Description

This command has no keywords or arguments.

Command Default

LISP configuration is disabled.

Command Modes

global configuration

Command History

Release	Modification
Release 4.3.0	This command was introduced.
Release 4.3.1	Support for LISP *process-number* was added.

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
lisp	read, write

This example shows how to configure LISP configuration mode:

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#router lisp
RP/0/RSP0/CPU0:router(config-lisp)#
```

Related Commands

Command	Description
address-family (LISP), on page 1949	Enters Locator ID and separation protocol (LISP) address family configuration mode.
itr map-resolver, on page 1965	Configures an IPv4 or IPv6 locator address of the LISP Map-Resolver to be used by the ITR.

Command	Description
loc-reach-algorithm rloc-probing, on page 1972	Configures locator reachability algorithm, RLOC Probing, and determines reachability status for other iBGP peers through the IGP domain.
map-cache-limit, on page 1974	Configures the maximum limit of IPv4 LISP or IPv6 LISP map-cache entries allowed to be stored by the router.
map-cache, on page 1975	Configures a static IPv4 EID-to-RLOC or static IPv6 EID-to-RLOC mapping relationship and its associated traffic policy, or statically configures the packet handling behavior associated with a destination IPv4 EID-prefix or a destination IPv6 EID-prefix.
proxy-etr, on page 1984	Configures a router to act as an IPv4 or IPv6 LISP Proxy Egress Tunnel Router (PETR).
proxy-itr, on page 1986	Configures a router to act as an IPv4 or IPv6 LISP Proxy Ingress Tunnel Router (PITR).

show lisp decapsulation filter

To display source Routing Locator (RLOC) addresses for specified parameters and the corresponding RLOC **show lisp decapsulation filter** command in privileged EXEC mode.

show lisp decapsulation filter [*IPv4-rloc-addressIPv4-rloc-address*] [**eid-table** *eid-table-vrf*|**instance-id** *iid*]

Syntax Description

IPv4-rloc-address	(Optional) Source RLOC address. If you want to know how a specific IPv4 RLOC address was configured, use this option.
IPv6-rloc-address	(Optional) Source RLOC address. If you want to know how a specific IPv6 RLOC address was configured, use this option.
eid-table *eid-table-vrf*	(Optional) Specifies the EID table and the associated VRF. Source RLOC addresses corresponding to the VRF is displayed.
instance-id *iid*	(Optional) Specifies the instance ID. Source RLOC addresses corresponding to the specified instance ID is displayed.

Command Modes

EXEC mode

Command History

Release	Modification
Release 5.3.0	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

Task ID

Task ID	Operation
sysmgr	read

The following sample output from the **show lisp decapsulation filter** command displays source RLOC address configuration details for a specific EID Instance ID. The RLOC address configuration details (whether it is manually configured or discovered) on a (P)xTR is displayed in the above table.

```
RP/0/RSP0/CPU0:router#show lisp decapsulation filter instance-id 0

LISP decapsulation filter for EID-table default (IID 0), 3 entries
Source RLOC   Added by
10.0.0.1      Config
10.0.0.5      209.165.200.230 209.165.200.232
10.0.0.6      Config 209.165.200.230
```

show lisp session

To display a current list of reliable transport (TCP) sessions, use the **show lisp session** command in privileged EXEC mode.

show lisp [**vrf vrf-name**] **session** [**established** | **peer-address**]

Syntax Description

vrf *vrf-name*	(Optional) Specifies the VRF instance. The transport session information for this VRF instance will be displayed.
session	(Optional) Specifies that reliable transport session information is displayed. If there are multiple transport sessions due to multiple roles, you can view information for all the sessions.
established	(Optional) Displays transport session information for established connections.
peer-address	(Optional) Specifies the peer IP address.

Command Modes

EXEC mode

Command History

Release	Modification
Release 5.3.0	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

Task ID

Task ID	Operation
sysmgr	read

The following sample output from the **show lisp session** command displays transport session information for a LISP VRF instance:

```
RP/0/RSP0/CPU0:router#show lisp session

Sessions for VRF default, total: 8, established: 7
Peer               State            Up/Down        In/Out     Users
2001:DB8:A:1::2    Up               00:04:13        2/7        2
2001:DB8:A:2::2    Up               00:04:13        2/7        2
2001:DB8:A:3::2    Up               00:03:53        2/7        2
2001:DB8:B:1::2    Up               00:04:04        2/6        2
2001:DB8:B:2::2    Init             never           0/0        1
2001:DB8:C:1::2    Up               00:03:55        2/6        2
2001:DB8:C:2::2    Up               00:03:54        2/6        2
2001:DB8:E:F::2    Up               00:04:04        6/19       4
```

show lisp site rloc members

To display source Routing Locator (RLOC) addresses for specified parameters and the corresponding RLOC **show lisp site rloc members** command in privileged EXEC mode.

show lisp [**instance-id** *iid*] [**site rloc members registrations** *rloc-address*]

Syntax Description

instance-id *iid*	(Optional) Specifies the instance ID for which the RLOC addresses will be displayed.
registrations	((Optional) Specifies that RLOC EID instance membership registration details be displayed.
rloc-address	(Optional) IPv4 or IPv6 RLOC address. If you want to view details for a specific RLOC address, you need to use this option.

Command Modes

EXEC mode

Command History

Release	Modification
Release 5.3.0	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

Task ID

Task ID	Operation
sysmgr	read

The following sample output from the show lisp decapsulation filter command displays source RLOC address configuration details for a specific EID Instance ID. The RLOC address configuration details (whether it is manually configured or discovered) on a MS/MR is displayed in the above table.

```
RP/0/RSP0/CPU0:router#show lisp site rloc members

LISP RLOC membership for EID table default (IID 0), 2 entries
RLOC          Origin                  Valid
10.0.1.2      registration            Yes
10.0.2.2      config & registration   Yes
```

The Origin column displays configuration details of the RLOC member. If an RLOC address is manually configured, automatically gleaned from received registrations, or both, the details are displayed. The Valid column shows whether the RLOC is a valid member that is distributed to (P)xTRs. A listed RLOC may not be valid if it is gleaned from registrations but the “override” option is used in the “modify-discovered” configuration and the specified locator-set does not include the RLOC.

show lisp site

To display configured LISP sites on a Locator/ID Separation Protocol (LISP) map server, use the **show lisp site** command in privileged EXEC mode.

show lisp site [**IPv4-dest-EID** | **IPv4-dest-EID-prefix** | **IPv6-dest-EID** | **IPv6-dest-EID-prefix**] | [**name** *site-name*] | [**detail**]

Syntax Description

IPv4-dest-EID	(Optional) Displays LISP site information matching this destination endpoint identifier (EID).
IPv4-dest-EID-prefix	(Optional) Displays LISP site information matching this destination EID prefix.
IPv6-dest-EID	(Optional) Displays LISP site information matching this destination EID.
IPv6-dest-EID-prefix	(Optional) Displays LISP site information matching this destination EID prefix.
name*site-name*	(Optional) Displays LISP site information matching this site name.
detail	(Optional) Increases the detail of all displayed LISP site information when no other parameters are used.

Command Modes

EXEC mode

Command History

Release	Modification
Release 5.3.0	This command was introduced.

Usage Guidelines

This command is used on a LISP map server to display information related to configured LISP sites. The displayed output indicates, among other things, whether a site is actively registered. When the base form of the command is used (**show lisp site**), summary information related to all configured LISP sites is displayed. When the IPv4-dest-EID form is used, a longest match is done to return the site with the best matching EID prefix and the displayed information applies specifically to that LISP site. When the IPv4-dest-EID-prefix form is used, an exact match is done to return the site configured with the EID prefix and the displayed information applies specifically to that LISP site. When the site-name form is used, the displayed information contains all EID prefixes configured for the named LISP site. When the **detail** keyword is added, all available details for the specific command form are presented.

Task ID

Task ID	Operation
sysmgr	read

The following sample output from the **show lisp site** command displays summary information related to all configured LISP sites:

```
RP/0/RSP0/CPU0:router#show lisp site

LISP Site Registration Information
Site Name       Last      Up          Who Last         EID Prefix
```

```
                Register                Registered
site1-xtr       00:00:04  yes           10.0.2.1          192.168.1.0/24
                00:00:04  yes           10.0.2.1          2001:DB8:A::/48
site2-xtr       00:00:35  yes           10.0.9.1          192.168.11.0/24
                00:00:35  yes           10.0.10.1         2001:DB8:B::/48
```

Related Commands

Command	Description
address-family (LISP), on page 1949	Enters Locator ID and separation protocol (LISP) address family configuration mode.
itr map-resolver, on page 1965	Configures an IPv4 or IPv6 locator address of the LISP Map-Resolver to be used by the ITR.
loc-reach-algorithm rloc-probing, on page 1972	Configures locator reachability algorithm, RLOC Probing, and determines reachability status for other iBGP peers through the IGP domain.
map-cache-limit, on page 1974	Configures the maximum limit of IPv4 LISP or IPv6 LISP map-cache entries allowed to be stored by the router.
map-cache, on page 1975	Configures a static IPv4 EID-to-RLOC or static IPv6 EID-to-RLOC mapping relationship and its associated traffic policy, or statically configures the packet handling behavior associated with a destination IPv4 EID-prefix or a destination IPv6 EID-prefix.
proxy-etr, on page 1984	Configures a router to act as an IPv4 or IPv6 LISP Proxy Egress Tunnel Router (PETR).
proxy-itr, on page 1986	Configures a router to act as an IPv4 or IPv6 LISP Proxy Ingress Tunnel Router (PITR).

solicit-map-request

To configure the solicit map request (SMR) handling, use the **solicit-map-request** command in the LISP address family configuration mode. To disable solicit map request handling, use the **no** form of this command.

solicit-map-request {**ignore** | **max-per-entry** *number* | **suppression-time** *seconds*}
no solicit-map-request {**ignore** | **max-per-entry** *number* | **suppression-time** *seconds*}

Syntax Description

ignore	Ignores an IPv4 or IPv6 map-request message that has the solicit-map-request (SMR) bit set.
max-per-entry	Specifies the maximum number of solicit-map-requests (SMRs) for addresses under a map-cache entry.
number	Specifies the maximum number of SMRs. The value ranges from 1 to 100.
suppression-time	Specifies how long to suppress repeated solicit-map-requests (SMRs) for the same address.
seconds	Specifies the seconds to suppress repeated SMRs for the same address. The value ranges from 1 to 600.

Command Default

Command Modes

LISP IPv4 address family

LISP IPv6 address family

Command History

Release	Modification
Release 4.3.1	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	Operation
lisp	read, write

This example shows how to configure the **solicit-map-request** command to ignore a map-request message, to specify a maximum number of 50 SMRs, and to specify a suppression time of 500 s:

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#router lisp
RP/0/RSP0/CPU0:router(config-lisp)#address-family ipv6 unicast
RP/0/RSP0/CPU0:router(config-lisp-af)#solicit-map-request ignore
```

```
RP/0/RSP0/CPU0:router(config-lisp-af)#solicit-map-request max-per-entry 50
RP/0/RSP0/CPU0:router(config-lisp-af)#solicit-map-request suppression-time 500
```

Related Commands

Command	Description
itr map-resolver, on page 1965	Configures an IPv4 or IPv6 locator address of the LISP Map-Resolver to be used by the ITR.
etr, on page 1958	Enables Egress Tunnel Router (ETR) functionality.
address-family (LISP), on page 1949	Enters Locator ID and separation protocol (LISP) address family configuration mode.
router lisp, on page 1989	Enters Locator and ID Separation Protocol (LISP) configuration mode.

use-petr

To specify that packets hitting a forward-native map-cache entry should be LISP encapsulated and forwarded to a PETR, instead of attempting to forward them natively, use the **use-petr** command in LISP address family configuration mode. To revert to the default behavior of forwarding packets natively, if they hit a forward-native map cache entry, use the **no** form of this command.

use-petr *ip_address* [**priority** *priority* **weight** *weight*]
no use-petr

Syntax Description

ip_address	Specifies the IPv4 or IPv6 locator address of the PETR.
priority	(Optional) Specifies the priority assigned to this PETR.
priority	Specifies the value of the priority assigned to this PETR. This value ranges from 0 to 255. A lower value indicates a higher priority.
weight	(Optional) Specifies the percentage of traffic to be load-shared.
weight	Specifies the weight in value of the percentage of traffic to be load-shared. The value ranges from 0 to 100.

Command Default

PETR services are disabled by default.

Command Modes

LISP IPv4 address family

LISP IPv6 address family

Command History

Release	Modification
Release 4.3.1	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.

Multiple PETRs can be configured. If you configure a priority and weight for one of them, you must configure a priority and weight for all of the PETRs.

Task ID

Task ID	Operation
lisp	read, write

This example shows how to use a single PETR:

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#router lisp
```

```
RP/0/RSP0/CPU0:router(config-lisp)#address-family ipv4 unicast
RP/0/RSP0/CPU0:router(config-lisp-af)#use-petr  2001:db8::1
```

This example shows how to configure a fallback PETR. This type of configuration is useful if, by default, you want to reach your PETR over IPv6, but in case you loose IPv6 connectivity to the outside world, you fall back to using IPv4. That way, your IPv6 EID prefix will maintain IPv6 connectivity to the outside world regardless of whether or not you have IPv6 connectivity at the ITR.

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#router lisp
RP/0/RSP0/CPU0:router(config-lisp)#address-family ipv6 unicast
RP/0/RSP0/CPU0:router(config-lisp-af)# use-petr 2001:db8::1 priority 1 weight 100
RP/0/RSP0/CPU0:router(config-lisp-af)#use-petr 64.10.10.1 priority 2 weight 100
```

Related Commands

Command	Description
itr map-resolver, on page 1965	Configures an IPv4 or IPv6 locator address of the LISP Map-Resolver to be used by the ITR.
address-family (LISP), on page 1949	Enters Locator ID and separation protocol (LISP) address family configuration mode.
router lisp, on page 1989	Enters Locator and ID Separation Protocol (LISP) configuration mode.