



Routing Command Reference for Cisco NCS 6000 Series Routers

First Published: 2013-09-20

Last Modified: 2017-09-01

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 NON-INFRINGEMENT 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)

© 2017–2018 Cisco Systems, Inc. All rights reserved.



CONTENTS

PREFACE

Preface xxvii

Changes to This Document xxvii

Communications, Services, and Additional Information xxvii

CHAPTER 1

BGP Commands 1

- address-family (BGP) 6
- advertisement-interval 8
- af-group 10
- aggregate-address 12
- aigp 14
- aigp send-cost-community 15
- allocate-label 17
- allow vpn default-originate 18
- allowas-in 19
- as-format 20
- as-path-loopcheck out disable 21
- attribute-filter group 22
- attribute large-community 23
- bfd (BGP) 24
- bgp as-path-loopcheck 26
- bgp attribute-download 27
- bgp auto-policy-soft-reset disable 29
- bgp bestpath as-path ignore 30
- bgp bestpath compare-routerid 32
- bgp bestpath cost-community ignore 33
- bgp bestpath med always 34

| | |
|-----------------------------------------|----|
| bgp bestpath med confed | 36 |
| bgp bestpath med missing-as-worst | 38 |
| bgp bestpath origin-as allow invalid | 39 |
| bgp bestpath origin-as use validity | 41 |
| bgp bestpath aigp ignore | 42 |
| bgp bestpath as-path multipath-relax | 43 |
| bgp client-to-client reflection disable | 44 |
| bgp cluster-id | 46 |
| bgp confederation identifier | 48 |
| bgp confederation peers | 50 |
| bgp dampening | 52 |
| bgp default local-preference | 54 |
| bgp enforce-first-as disable | 55 |
| bgp fast-external-falover disable | 57 |
| bgp graceful-restart | 58 |
| bgp graceful-restart graceful-reset | 60 |
| bgp graceful-restart purge-time | 61 |
| bgp graceful-restart restart-time | 62 |
| bgp graceful-restart stalepath-time | 63 |
| bgp label-delay | 65 |
| bgp log neighbor changes disable | 66 |
| bgp maximum neighbor | 68 |
| bgp nexthop resolution allow-default | 69 |
| bgp redistribute-internal | 70 |
| bgp router-id | 72 |
| bgp scan-time | 73 |
| bgp update-delay | 74 |
| capability additional-paths receive | 75 |
| capability additional-paths send | 76 |
| capability orf prefix | 77 |
| capability suppress | 79 |
| clear bgp | 82 |
| cef consistency-hashing auto-recovery | 84 |
| clear bgp dampening | 85 |

| | |
|------------------------------------------|-----|
| clear bgp external | 87 |
| clear bgp flap-statistics | 88 |
| clear bgp nexthop performance-statistics | 90 |
| clear bgp nexthop registration | 92 |
| clear bgp peer-drops | 94 |
| clear bgp performance-statistics | 95 |
| clear bgp rpki | 96 |
| clear bgp self-originated | 98 |
| clear bgp shutdown | 100 |
| clear bgp soft | 101 |
| cluster-id | 104 |
| cluster-id allow-equal | 106 |
| default-information originate (BGP) | 107 |
| default-metric (BGP) | 108 |
| default-originate | 109 |
| description (BGP) | 111 |
| distance bgp | 112 |
| dscp (BGP) | 114 |
| ebgp-multihop | 116 |
| enforce-first-as | 118 |
| enforce-first-as-disable | 120 |
| ibgp policy out enforce-modifications | 122 |
| ignore-connected-check | 123 |
| keychain | 124 |
| keychain-disable | 126 |
| keychain inheritance-disable | 128 |
| local-as | 130 |
| maximum-paths (BGP) | 132 |
| maximum-prefix (BGP) | 134 |
| mpls activate (BGP) | 137 |
| multipath | 138 |
| neighbor (BGP) | 139 |
| neighbor-group | 141 |
| network (BGP) | 143 |

| | |
|----------------------------------|-----|
| network backdoor | 145 |
| next-hop-self | 147 |
| next-hop-unchanged | 149 |
| nexthop route-policy | 151 |
| nexthop trigger-delay | 153 |
| nsr (BGP) | 155 |
| nsr disable (BGP) | 157 |
| orf | 159 |
| password (BGP) | 161 |
| password-disable | 163 |
| precedence | 165 |
| receive-buffer-size | 167 |
| redistribute (BGP) | 169 |
| remote-as (BGP) | 173 |
| remove-private-as | 176 |
| route-policy (BGP) | 178 |
| route-reflector-client | 180 |
| router bgp | 182 |
| send-buffer-size | 183 |
| send-community-ebgp | 185 |
| send-extended-community-ebgp | 187 |
| session-group | 189 |
| session-open-mode | 191 |
| show bgp | 193 |
| show bgp large-community | 203 |
| send-community-ebgp | 205 |
| show bgp update out | 207 |
| show bgp update in error process | 208 |
| show bgp update out filter-group | 209 |
| show bgp update out process | 210 |
| show bgp update out sub-group | 212 |
| show bgp update out update-group | 213 |
| show bgp advertised | 214 |
| show bgp af-group | 220 |

| | |
|--------------------------------|-----|
| show bgp attribute-key | 223 |
| show bgp cidr-only | 226 |
| show bgp community | 230 |
| show bgp convergence | 235 |
| show bgp dampened-paths | 238 |
| show bgp flap-statistics | 242 |
| show bgp inconsistent-as | 247 |
| show bgp labels | 251 |
| show bgp neighbor-group | 252 |
| show bgp neighbors | 256 |
| show bgp neighbors nsr | 273 |
| show bgp nexthops | 275 |
| show bgp nsr | 277 |
| show bgp paths | 281 |
| show bgp policy | 283 |
| show bgp process | 290 |
| show bgp regexp | 308 |
| show bgp route-policy | 312 |
| show bgp session-group | 316 |
| show bgp sessions | 319 |
| show bgp summary | 320 |
| show bgp summary nsr | 324 |
| show bgp table | 327 |
| show bgp truncated-communities | 329 |
| show bgp update-group | 333 |
| show protocols (BGP) | 336 |
| show tcp brief | 338 |
| shutdown (BGP) | 339 |
| socket receive-buffer-size | 341 |
| socket send-buffer-size | 343 |
| soft-reconfiguration inbound | 345 |
| speaker-id | 348 |
| table-policy | 349 |
| timers (BGP) | 351 |

- timers bgp 353
- ttl-security 355
- update limit 358
- update limit address-family 359
- update limit sub-group 360
- update in error-handling basic disable 362
- update in error-handling extended 363
- update in labeled-unicast equivalent 364
- update out logging 365
- update-source 366
- use 368
- weight 372

CHAPTER 2**BGP Flowspec Commands 375**

- class-map type traffic (BGP-flowspec) 376
- class type traffic 377
- drop (BGP-flowspec) 378
- flowspec 379
- flowspec disable 380
- local-install 381
- match destination-address 382
- match destination-port 383
- match dscp 384
- match fragment-type 387
- match icmp code 388
- match icmp type 389
- match packet length 390
- match protocol 391
- match source-address 393
- match source-port 394
- match tcp flag 395
- policy-map 396
- redirect (BGP Flowspec) 398
- service-policy 399

show flowspec 400

CHAPTER 3

BFD Commands 401

address-family ipv4 unicast (BFD) 402

bfd 404

bfd address-family ipv4 destination 406

bfd address-family ipv4 fast-detect 407

bfd address-family ipv4 minimum-interval 408

bfd address-family ipv4 multiplier 411

bfd address-family ipv4 timers 414

bundle coexistence bob-blb 416

bfd fast-detect 418

bfd minimum-interval 421

bfd multipath include location 425

bfd multiplier 427

clear bfd counters 429

echo disable 431

echo ipv4 source 433

echo latency detect 435

echo startup validate 437

interface (BFD) 439

ipv6 checksum 442

multihop ttl-drop-threshold 444

show bfd 446

show bfd client 449

show bfd counters 451

show bfd mib session 454

show bfd multipath 457

show bfd session 458

show bfd summary 463

CHAPTER 4

EIGRP Commands 465

address-family (EIGRP) 467

authentication keychain 468

| | |
|-------------------------------|-----|
| auto-summary (EIGRP) | 469 |
| bandwidth-percent (EIGRP) | 470 |
| clear eigrp neighbors | 471 |
| clear eigrp topology | 472 |
| default-information | 473 |
| default-metric (EIGRP) | 474 |
| distance (EIGRP) | 475 |
| hello-interval (EIGRP) | 477 |
| hold-time (EIGRP) | 478 |
| interface (EIGRP) | 480 |
| log-neighbor-changes | 482 |
| log-neighbor-warnings | 483 |
| maximum-paths (EIGRP) | 484 |
| metric (EIGRP) | 485 |
| metric maximum-hops | 487 |
| metric weights | 488 |
| neighbor (EIGRP) | 490 |
| next-hop-self disable | 492 |
| nsf disable (EIGRP) | 493 |
| passive-interface (EIGRP) | 494 |
| redistribute (EIGRP) | 495 |
| route-policy (EIGRP) | 497 |
| router eigrp | 498 |
| router-id (EIGRP) | 499 |
| show protocols (EIGRP) | 500 |
| split-horizon disable (EIGRP) | 503 |
| stub (EIGRP) | 504 |
| summary-address (EIGRP) | 506 |
| timers active-time | 508 |
| timers nsf route-hold (EIGRP) | 509 |
| variance | 510 |

CHAPTER 5**IS-IS Commands 511**

| | |
|------------------------|-----|
| address-family (IS-IS) | 514 |
|------------------------|-----|

| | |
|------------------------------------------------------|-----|
| adjacency-check disable | 516 |
| adjacency stagger | 518 |
| attached-bit receive ignore | 520 |
| attached-bit send | 521 |
| circuit-type | 523 |
| clear isis process | 525 |
| clear isis route | 526 |
| clear isis statistics | 528 |
| csnp-interval | 530 |
| default-information originate (IS-IS) | 532 |
| disable (IS-IS) | 534 |
| distance (IS-IS) | 535 |
| distribute-list in (IS-IS) | 537 |
| fast-reroute per-link (IS-IS) | 539 |
| fast-reroute per-prefix (IS-IS) | 541 |
| fast-reroute per-link priority-limit (IS-IS) | 543 |
| fast-reroute per-prefix load-sharing disable (IS-IS) | 544 |
| fast-reroute per-prefix tiebreaker (IS-IS) | 545 |
| hello-interval (IS-IS) | 547 |
| hello-multiplier | 549 |
| hello-padding | 551 |
| hello-password | 553 |
| hello-password keychain | 555 |
| hello-password accept | 557 |
| hostname dynamic disable | 559 |
| ignore-lsp-errors | 560 |
| instance-id | 561 |
| interface (IS-IS) | 562 |
| ipfrr lfa | 564 |
| ipfrr lfa exclude interface | 566 |
| is-type | 568 |
| link-down fast-detect | 570 |
| link-group | 571 |
| log adjacency changes (IS-IS) | 572 |

| | |
|-------------------------------------------------|-----|
| log pdu drops | 573 |
| lsp fast-flood threshold | 574 |
| lsp-gen-interval | 575 |
| lsp-interval | 577 |
| lsp-mtu | 578 |
| lsp-password | 580 |
| lsp-password accept | 582 |
| lsp-refresh-interval | 584 |
| maximum-paths (IS-IS) | 586 |
| maximum-redistributed-prefixes (IS-IS) | 587 |
| max-lsp-lifetime | 588 |
| mesh-group (IS-IS) | 589 |
| metric (IS-IS) | 591 |
| metric-style narrow | 593 |
| metric-style transition | 595 |
| metric-style wide | 597 |
| microloop avoidance | 599 |
| microloop avoidance segment-routing | 601 |
| min-lsp-arrivaltime | 603 |
| mpls ldp auto-config | 605 |
| mpls ldp sync (IS-IS) | 606 |
| mpls traffic-eng (IS-IS) | 607 |
| mpls traffic-eng path-selection ignore overload | 609 |
| mpls traffic-eng router-id (IS-IS) | 611 |
| mpls traffic-eng srlg (IS-IS) | 613 |
| net | 615 |
| nsf (IS-IS) | 617 |
| nsf interface-expires | 619 |
| nsf interface-timer | 621 |
| nsf lifetime (IS-IS) | 623 |
| passive (IS-IS) | 624 |
| point-to-point | 625 |
| priority (IS-IS) | 626 |
| propagate level | 628 |

| | |
|-------------------------------------------|-----|
| redistribute (IS-IS) | 630 |
| retransmit-interval (IS-IS) | 634 |
| retransmit-throttle-interval | 636 |
| router isis | 638 |
| route source first-hop | 639 |
| set-overload-bit | 640 |
| set-attached-bit | 642 |
| show isis | 644 |
| show isis adjacency | 646 |
| show isis adjacency-log | 648 |
| show isis checkpoint adjacency | 650 |
| show isis checkpoint interface | 652 |
| show isis checkpoint lsp | 654 |
| show isis database | 656 |
| show isis database-log | 658 |
| show isis fast-reroute | 660 |
| show isis hostname | 663 |
| show isis interface | 665 |
| show isis lsp-log | 669 |
| show isis mesh-group | 671 |
| show isis mpls traffic-eng adjacency-log | 673 |
| show isis mpls traffic-eng advertisements | 675 |
| show isis mpls traffic-eng tunnel | 677 |
| show isis neighbors | 679 |
| show isis protocol | 682 |
| show isis route | 684 |
| show isis spf-log | 687 |
| show isis statistics | 694 |
| show isis topology | 698 |
| show protocols (IS-IS) | 701 |
| shutdown (IS-IS) | 704 |
| single-topology | 705 |
| snmp-server traps isis | 706 |
| spf-interval | 707 |

spf prefix-priority (IS-IS) 709
summary-prefix (IS-IS) 711
suppressed 713
tag (IS-IS) 714
trace (IS-IS) 715

CHAPTER 6**OSPF Commands 717**

address-family (OSPF) 720
adjacency stagger 721
area (OSPF) 723
authentication (OSPF) 725
authentication-key (OSPF) 728
auto-cost (OSPF) 730
capability opaque disable 732
clear ospf process 733
clear ospf redistribution 735
clear ospf routes 736
clear ospf statistics 737
clear ospf statistics interface 739
cost (OSPF) 740
cost-fallback (OSPF) 742
database-filter all out (OSPF) 744
dead-interval (OSPF) 745
default-cost (OSPF) 747
default-information originate (OSPF) 749
default-metric (OSPF) 751
demand-circuit (OSPF) 753
distance (OSPF) 755
distance ospf 758
distribute-list 760
fast-reroute (OSPFv2) 762
fast-reroute per-link exclude interface 764
fast-reroute per-prefix exclude interface (OSPFv2) 766
fast-reroute per-prefix lfa-candidate (OSPFv2) 767

| | |
|-----------------------------------------------------|-----|
| fast-reroute per-prefix remote-lfa (OSPFv2) | 768 |
| fast-reroute per-prefix use-candidate-only (OSPFv2) | 770 |
| flood-reduction (OSPF) | 771 |
| hello-interval (OSPF) | 773 |
| interface (OSPF) | 775 |
| log adjacency changes (OSPF) | 777 |
| loopback stub-network | 778 |
| max-lsa | 779 |
| max-metric | 781 |
| maximum interfaces (OSPF) | 784 |
| maximum paths (OSPF) | 785 |
| maximum redistributed-prefixes (OSPF) | 787 |
| message-digest-key | 789 |
| mpls ldp auto-config (OSPF) | 792 |
| mpls ldp sync (OSPF) | 793 |
| mpls ldp sync-igp-shortcuts | 794 |
| mpls traffic-eng ldp-sync-update (OSPF) | 796 |
| mtu-ignore (OSPF) | 797 |
| multi-area-interface | 799 |
| neighbor (OSPF) | 801 |
| neighbor database-filter all out | 803 |
| network (OSPF) | 804 |
| nsf (OSPF) | 806 |
| nsf flush-delay-time (OSPF) | 808 |
| nsf interval (OSPF) | 809 |
| nsf lifetime (OSPF) | 810 |
| nssa (OSPF) | 811 |
| ospf name-lookup | 813 |
| packet-size (OSPF) | 814 |
| passive (OSPF) | 816 |
| priority (OSPF) | 818 |
| protocol shutdown | 820 |
| queue dispatch flush-lsa | 821 |
| queue dispatch incoming | 823 |

| | |
|---------------------------------|-----|
| queue dispatch rate-limited-lsa | 825 |
| queue dispatch spf-lsa-limit | 827 |
| queue limit | 828 |
| range (OSPF) | 830 |
| redistribute (OSPF) | 832 |
| retransmit-interval (OSPF) | 837 |
| route-policy (OSPF) | 839 |
| router-id (OSPF) | 840 |
| router ospf | 842 |
| security ttl (OSPF) | 844 |
| show ospf | 846 |
| show ospf border-routers | 849 |
| show ospf flood-list | 851 |
| show ospf interface | 853 |
| show ospf mpls traffic-eng | 855 |
| show ospf message-queue | 860 |
| show ospf neighbor | 863 |
| show ospf request-list | 870 |
| show ospf retransmission-list | 872 |
| show ospf routes | 874 |
| show ospf sham-links | 878 |
| show ospf statistics interface | 879 |
| show ospf summary-prefix | 881 |
| show ospf virtual-links | 883 |
| show protocols (OSPF) | 885 |
| snmp trap rate-limit (OSPF) | 888 |
| spf prefix-priority (OSPFv2) | 889 |
| stub (OSPF) | 891 |
| summary-prefix (OSPF) | 893 |
| timers lsa group-pacing | 895 |
| timers lsa min-arrival | 896 |
| timers lsa refresh | 897 |
| timers throttle lsa all (OSPF) | 899 |
| timers throttle spf (OSPF) | 902 |

transmit-delay (OSPF) 904
ucmp (OSPFv2) 906
ucmp delay-interval (OSPFv2) 908
ucmp exclude interface (OSPFv2) 909
virtual-link (OSPF) 911

CHAPTER 7
OSPFv3 Commands 913

address-family (OSPFv3) 915
area (OSPFv3) 916
authentication (OSPFv3) 917
auto-cost (OSPFv3) 919
clear ospfv3 process 921
clear ospfv3 redistribution 923
clear ospfv3 routes 924
clear ospfv3 statistics 925
cost (OSPFv3) 927
database-filter all out (OSPFv3) 929
dead-interval (OSPFv3) 931
default-cost (OSPFv3) 933
default-information originate (OSPFv3) 935
default-metric (OSPFv3) 937
demand-circuit (OSPFv3) 939
distance ospfv3 941
distribute-list prefix-list in 943
distribute-list prefix-list out 945
encryption 947
flood-reduction (OSPFv3) 949
graceful-restart (OSPFv3) 951
hello-interval (OSPFv3) 952
instance (OSPFv3) 954
interface (OSPFv3) 956
log adjacency changes (OSPFv3) 958
maximum interfaces (OSPFv3) 959
maximum paths (OSPFv3) 960

| | |
|-----------------------------------------|------|
| maximum redistributed-prefixes (OSPFv3) | 961 |
| mtu-ignore (OSPFv3) | 963 |
| neighbor (OSPFv3) | 965 |
| network (OSPFv3) | 967 |
| nssa (OSPFv3) | 969 |
| nsr (OSPFv3) | 971 |
| ospfv3 name-lookup | 973 |
| packet-size (OSPFv3) | 974 |
| passive (OSPFv3) | 975 |
| priority (OSPFv3) | 977 |
| protocol shutdown (OSPFv3) | 979 |
| range (OSPFv3) | 980 |
| redistribute (OSPFv3) | 982 |
| retransmit-interval (OSPFv3) | 986 |
| router-id (OSPFv3) | 988 |
| router ospfv3 | 990 |
| show ospfv3 | 991 |
| show ospfv3 border-routers | 995 |
| show ospfv3 database | 997 |
| show ospfv3 flood-list | 1009 |
| show ospfv3 interface | 1011 |
| show ospfv3 message-queue | 1014 |
| show ospfv3 neighbor | 1016 |
| show ospfv3 request-list | 1022 |
| show ospfv3 retransmission-list | 1024 |
| show ospfv3 routes | 1026 |
| show ospfv3 summary-prefix | 1028 |
| show ospfv3 virtual-links | 1030 |
| show protocols (OSPFv3) | 1032 |
| stub (OSPFv3) | 1034 |
| stub-router | 1036 |
| summary-prefix (OSPFv3) | 1038 |
| timers lsa arrival | 1040 |
| timers pacing flood | 1041 |

| | |
|----------------------------------|------|
| timers pacing lsa-group | 1043 |
| timers pacing retransmission | 1045 |
| timers throttle lsa all (OSPFv3) | 1047 |
| timers throttle spf (OSPFv3) | 1049 |
| trace (OSPFv3) | 1051 |
| transmit-delay (OSPFv3) | 1053 |
| virtual-link (OSPFv3) | 1055 |

CHAPTER 8
RIB Commands 1057

| | |
|-------------------------------------------|------|
| address-family next-hop dampening disable | 1059 |
| clear route | 1060 |
| lcc | 1062 |
| rcc | 1063 |
| recursion-depth-max | 1065 |
| router rib | 1066 |
| rump always-replicate | 1067 |
| show lcc statistics | 1068 |
| show rcc | 1070 |
| show rcc statistics | 1072 |
| show rib | 1074 |
| show rib afi-all | 1076 |
| show rib attributes | 1078 |
| show rib client-id | 1079 |
| show rib clients | 1081 |
| show rib extcomms | 1082 |
| show rib firsthop | 1084 |
| show rib history | 1086 |
| show rib next-hop | 1088 |
| show rib opaques | 1090 |
| show rib protocols | 1092 |
| show rib recursion-depth-max | 1094 |
| show rib statistics | 1096 |
| show rib tables | 1098 |
| show rib trace | 1099 |

show route 1101
show route backup 1106
show route best-local 1109
show route connected 1111
show route local 1113
show route longer-prefixes 1115
show route next-hop 1117
show route quarantined 1119
show route resolving-next-hop 1121
show route static 1123
show route summary 1125

CHAPTER 9**RIP Commands 1129**

authentication keychain mode (RIP) 1130
auto-summary (RIP) 1132
broadcast-for-v2 1133
default-information originate (RIP) 1134
default-metric (RIP) 1135
distance (RIP) 1136
interface (RIP) 1138
maximum-paths (RIP) 1140
metric-zero-accept 1141
neighbor (RIP) 1142
nsf (RIP) 1143
output-delay 1144
passive-interface (RIP) 1145
poison-reverse 1146
receive version 1148
redistribute (RIP) 1149
router rip 1152
route-policy (RIP) 1154
send version 1155
show protocols (RIP) 1156
show rip 1157

| | |
|--------------------------------|------|
| show rip database | 1159 |
| show rip interface | 1161 |
| show rip statistics | 1164 |
| site-of-origin (RIP) | 1166 |
| split-horizon disable (RIP) | 1168 |
| timers basic | 1170 |
| validate-update-source disable | 1172 |

CHAPTER 10 **Routing Policy Language Commands** 1173

| | |
|--------------------------|------|
| abort (RPL) | 1177 |
| add | 1179 |
| apply | 1181 |
| as-path in | 1183 |
| as-path is-local | 1185 |
| as-path length | 1187 |
| as-path neighbor-is | 1189 |
| as-path originates-from | 1192 |
| as-path passes-through | 1194 |
| as-path-set | 1196 |
| as-path unique-length | 1198 |
| community is-empty | 1200 |
| community matches-any | 1201 |
| community matches-every | 1203 |
| community matches-within | 1205 |
| community-set | 1207 |
| delete community | 1210 |
| delete large-community | 1212 |
| delete extcommunity rt | 1214 |
| destination in | 1216 |
| done | 1218 |
| drop | 1220 |
| edit | 1222 |
| end-global | 1225 |
| end-policy | 1226 |

end-set 1227

extcommunity rt is-empty 1229

extcommunity rt matches-any 1230

extcommunity rt matches-every 1232

extcommunity rt matches-within 1234

extcommunity-set cost 1236

extcommunity-set rt 1238

extcommunity-set soo 1240

extcommunity soo is-empty 1242

extcommunity soo matches-any 1243

extcommunity soo matches-every 1245

if 1247

local-preference 1253

large-community is-empty 1254

large-community matches-any 1256

large-community matches-every 1258

large-community matches-within 1260

large-community-set 1262

med 1263

next-hop in 1264

orf prefix in 1266

origin is 1268

ospf-area-set 1270

pass 1272

path-type is 1274

policy-global 1275

prefix-set 1277

prepend as-path 1279

protocol 1281

rd in 1283

replace as-path 1284

rib-has-route 1286

route-has-label 1288

route-policy (RPL) 1289

| | |
|-------------------------------|------|
| route-type is | 1291 |
| rpl editor | 1293 |
| rpl maximum | 1294 |
| set aigp-metric | 1296 |
| set community | 1297 |
| set dampening | 1299 |
| set eigrp-metric | 1301 |
| set extcommunity cost | 1303 |
| set extcommunity rt | 1305 |
| set ip-precedence | 1307 |
| set isis-metric | 1309 |
| set label | 1310 |
| set large-community | 1311 |
| set level | 1313 |
| set local-preference | 1314 |
| set med | 1315 |
| set metric-type (IS-IS) | 1317 |
| set metric-type (OSPF) | 1318 |
| set next-hop | 1319 |
| set origin | 1321 |
| set ospf-metric | 1322 |
| set path-selection | 1323 |
| set qos-group (RPL) | 1325 |
| set rib-metric | 1326 |
| set rip-metric | 1327 |
| set rip-tag | 1328 |
| set rtset route-limit | 1329 |
| set spf-priority | 1330 |
| set tag | 1331 |
| set traffic-index | 1332 |
| set weight | 1334 |
| show rpl | 1336 |
| show rpl active as-path-set | 1338 |
| show rpl active community-set | 1341 |

| | |
|-------------------------------------|------|
| show rpl active extcommunity-set | 1344 |
| show rpl active prefix-set | 1347 |
| show rpl active rd-set | 1350 |
| show rpl active route-policy | 1352 |
| show rpl as-path-set | 1354 |
| show rpl as-path-set attachpoints | 1356 |
| show rpl as-path-set references | 1359 |
| show rpl community-set | 1361 |
| show rpl community-set attachpoints | 1363 |
| show rpl community-set references | 1365 |
| show rpl extcommunity-set | 1368 |
| show rpl inactive as-path-set | 1371 |
| show rpl inactive community-set | 1374 |
| show rpl inactive extcommunity-set | 1377 |
| show rpl inactive prefix-set | 1380 |
| show rpl inactive rd-set | 1383 |
| show rpl inactive route-policy | 1385 |
| show rpl maximum | 1388 |
| show rpl policy-global references | 1390 |
| show rpl prefix-set | 1392 |
| show rpl prefix-set attachpoints | 1394 |
| show rpl prefix-set references | 1397 |
| show rpl rd-set | 1399 |
| show rpl rd-set references | 1401 |
| show rpl route-policy | 1403 |
| show rpl route-policy attachpoints | 1406 |
| show rpl route-policy inline | 1408 |
| show rpl route-policy references | 1410 |
| show rpl route-policy uses | 1413 |
| show rpl unused as-path-set | 1416 |
| show rpl unused community-set | 1419 |
| show rpl unused extcommunity-set | 1422 |
| show rpl unused prefix-set | 1424 |
| show rpl unused rd-set | 1427 |

| | |
|------------------------------|------|
| show rpl unused route-policy | 1429 |
| source in | 1432 |
| source rt | 1434 |
| suppress-route | 1435 |
| tag | 1436 |
| tag in | 1437 |
| tag-set | 1439 |
| unsuppress-route | 1440 |

| | | |
|-------------------|--------------------------------|-------------|
| CHAPTER 11 | Static Routing Commands | 1443 |
| | address-family (static) | 1444 |
| | maximum path (static) | 1445 |
| | route (static) | 1447 |
| | router static | 1450 |

| | | |
|-------------------|-----------------------------|-------------|
| CHAPTER 12 | RCMD Commands | 1451 |
| | router-convergence | 1452 |
| | monitor-convergence (IS-IS) | 1453 |
| | monitor-convergence (OSPF) | 1454 |
| | collect-diagnostics (RCMD) | 1455 |
| | event-buffer-size (RCMD) | 1457 |
| | max-events-stored (RCMD) | 1458 |
| | monitoring-interval (RCMD) | 1459 |
| | node disable (RCMD) | 1460 |
| | priority (RCMD) | 1462 |
| | protocol (RCMD) | 1464 |
| | storage-location | 1465 |



Preface

The *Routing Command Reference for Cisco NCS 6000 Series Routers* preface contains these sections:

- [Changes to This Document, on page xxvii](#)
- [Communications, Services, and Additional Information, on page xxvii](#)

Changes to This Document

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

Table 1: Changes to This Document

| Date | Change Summary |
|----------------|--------------------------------------------------------------------|
| September 2013 | Initial release of this document. |
| January 2014 | Republished with documentation updates for Release 5.0.1 features. |
| August 2014 | Republished with documentation updates for Release 5.2.1 features. |
| January 2015 | Republished with documentation updates for Release 5.2.3 features. |
| November 2016 | Republished with documentation updates for Release 6.1.2 features. |
| July 2017 | Republished with documentation updates for Release 6.2.2 features. |
| September 2017 | Republished with documentation updates for Release 6.3.1 features. |

Communications, Services, and Additional Information

- To receive timely, relevant information from Cisco, sign up at [Cisco Profile Manager](#).

- To get the business impact 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) IP Version 4 (IPv4), IP Version 6 (IPv6),) routing sessions.

For detailed information about BGP concepts, configuration tasks, and examples, see the *Implementing BGP* chapter in the *Routing Configuration Guide for Cisco NCS 6000 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.

- [address-family \(BGP\), on page 6](#)
- [advertisement-interval, on page 8](#)
- [af-group, on page 10](#)
- [aggregate-address, on page 12](#)
- [aigp, on page 14](#)
- [aigp send-cost-community, on page 15](#)
- [allocate-label, on page 17](#)
- [allow vpn default-originate, on page 18](#)
- [allowas-in, on page 19](#)
- [as-format, on page 20](#)
- [as-path-loopcheck out disable, on page 21](#)
- [attribute-filter group, on page 22](#)
- [attribute large-community, on page 23](#)
- [bfd \(BGP\), on page 24](#)
- [bgp as-path-loopcheck, on page 26](#)
- [bgp attribute-download, on page 27](#)
- [bgp auto-policy-soft-reset disable, on page 29](#)
- [bgp bestpath as-path ignore, on page 30](#)
- [bgp bestpath compare-routerid, on page 32](#)
- [bgp bestpath cost-community ignore, on page 33](#)
- [bgp bestpath med always, on page 34](#)
- [bgp bestpath med confed, on page 36](#)
- [bgp bestpath med missing-as-worst, on page 38](#)

- [bgp bestpath origin-as allow invalid](#), on page 39
- [bgp bestpath origin-as use validity](#), on page 41
- [bgp bestpath aigp ignore](#), on page 42
- [bgp bestpath as-path multipath-relax](#) , on page 43
- [bgp client-to-client reflection disable](#), on page 44
- [bgp cluster-id](#), on page 46
- [bgp confederation identifier](#), on page 48
- [bgp confederation peers](#), on page 50
- [bgp dampening](#), on page 52
- [bgp default local-preference](#), on page 54
- [bgp enforce-first-as disable](#), on page 55
- [bgp fast-external-fallover disable](#), on page 57
- [bgp graceful-restart](#), on page 58
- [bgp graceful-restart graceful-reset](#), on page 60
- [bgp graceful-restart purge-time](#), on page 61
- [bgp graceful-restart restart-time](#), on page 62
- [bgp graceful-restart stalepath-time](#), on page 63
- [bgp label-delay](#), on page 65
- [bgp log neighbor changes disable](#), on page 66
- [bgp maximum neighbor](#), on page 68
- [bgp nexthop resolution allow-default](#), on page 69
- [bgp redistribute-internal](#), on page 70
- [bgp router-id](#), on page 72
- [bgp scan-time](#), on page 73
- [bgp update-delay](#), on page 74
- [capability additional-paths receive](#), on page 75
- [capability additional-paths send](#), on page 76
- [capability orf prefix](#), on page 77
- [capability suppress](#) , on page 79
- [clear bgp](#), on page 82
- [cef consistency-hashing auto-recovery](#), on page 84
- [clear bgp dampening](#), on page 85
- [clear bgp external](#), on page 87
- [clear bgp flap-statistics](#), on page 88
- [clear bgp nexthop performance-statistics](#), on page 90
- [clear bgp nexthop registration](#), on page 92
- [clear bgp peer-drops](#), on page 94
- [clear bgp performance-statistics](#), on page 95
- [clear bgp rpki](#) , on page 96
- [clear bgp self-originated](#), on page 98
- [clear bgp shutdown](#), on page 100
- [clear bgp soft](#), on page 101
- [cluster-id](#), on page 104
- [cluster-id allow-equal](#) , on page 106
- [default-information originate \(BGP\)](#), on page 107
- [default-metric \(BGP\)](#), on page 108

- [default-originate](#), on page 109
- [description \(BGP\)](#), on page 111
- [distance bgp](#), on page 112
- [dscp \(BGP\)](#), on page 114
- [ebgp-multihop](#), on page 116
- [enforce-first-as](#), on page 118
- [enforce-first-as-disable](#), on page 120
- [ibgp policy out enforce-modifications](#), on page 122
- [ignore-connected-check](#), on page 123
- [keychain](#), on page 124
- [keychain-disable](#), on page 126
- [keychain inheritance-disable](#), on page 128
- [local-as](#), on page 130
- [maximum-paths \(BGP\)](#), on page 132
- [maximum-prefix \(BGP\)](#), on page 134
- [mpls activate \(BGP\)](#), on page 137
- [multipath](#), on page 138
- [neighbor \(BGP\)](#), on page 139
- [neighbor-group](#), on page 141
- [network \(BGP\)](#), on page 143
- [network backdoor](#), on page 145
- [next-hop-self](#), on page 147
- [next-hop-unchanged](#), on page 149
- [nexthop route-policy](#), on page 151
- [nexthop trigger-delay](#), on page 153
- [nsr \(BGP\)](#), on page 155
- [nsr disable \(BGP\)](#), on page 157
- [orf](#), on page 159
- [password \(BGP\)](#), on page 161
- [password-disable](#), on page 163
- [precedence](#), on page 165
- [receive-buffer-size](#), on page 167
- [redistribute \(BGP\)](#), on page 169
- [remote-as \(BGP\)](#), on page 173
- [remove-private-as](#), on page 176
- [route-policy \(BGP\)](#), on page 178
- [route-reflector-client](#), on page 180
- [router bgp](#), on page 182
- [send-buffer-size](#), on page 183
- [send-community-ebgp](#), on page 185
- [send-extended-community-ebgp](#), on page 187
- [session-group](#), on page 189
- [session-open-mode](#), on page 191
- [show bgp](#), on page 193
- [show bgp large-community](#), on page 203
- [send-community-ebgp](#), on page 205

- [show bgp update out](#), on page 207
- [show bgp update in error process](#), on page 208
- [show bgp update out filter-group](#), on page 209
- [show bgp update out process](#), on page 210
- [show bgp update out sub-group](#), on page 212
- [show bgp update out update-group](#), on page 213
- [show bgp advertised](#), on page 214
- [show bgp af-group](#), on page 220
- [show bgp attribute-key](#), on page 223
- [show bgp cidr-only](#), on page 226
- [show bgp community](#), on page 230
- [show bgp convergence](#), on page 235
- [show bgp dampened-paths](#), on page 238
- [show bgp flap-statistics](#), on page 242
- [show bgp inconsistent-as](#), on page 247
- [show bgp labels](#), on page 251
- [show bgp neighbor-group](#), on page 252
- [show bgp neighbors](#), on page 256
- [show bgp neighbors nsr](#), on page 273
- [show bgp nexthops](#), on page 275
- [show bgp nsr](#), on page 277
- [show bgp paths](#), on page 281
- [show bgp policy](#), on page 283
- [show bgp process](#), on page 290
- [show bgp regexp](#), on page 308
- [show bgp route-policy](#), on page 312
- [show bgp session-group](#), on page 316
- [show bgp sessions](#), on page 319
- [show bgp summary](#), on page 320
- [show bgp summary nsr](#), on page 324
- [show bgp table](#), on page 327
- [show bgp truncated-communities](#), on page 329
- [show bgp update-group](#), on page 333
- [show protocols \(BGP\)](#), on page 336
- [show tcp brief](#), on page 338
- [shutdown \(BGP\)](#), on page 339
- [socket receive-buffer-size](#), on page 341
- [socket send-buffer-size](#), on page 343
- [soft-reconfiguration inbound](#), on page 345
- [speaker-id](#), on page 348
- [table-policy](#), on page 349
- [timers \(BGP\)](#), on page 351
- [timers bgp](#), on page 353
- [ttl-security](#), on page 355
- [update limit](#), on page 358
- [update limit address-family](#), on page 359

- [update limit sub-group, on page 360](#)
- [update in error-handling basic disable, on page 362](#)
- [update in error-handling extended, on page 363](#)
- [update in labeled-unicast equivalent, on page 364](#)
- [update out logging, on page 365](#)
- [update-source, on page 366](#)
- [use, on page 368](#)
- [weight, on page 372](#)

address-family (BGP)

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

```
address-family {ipv4 {labeled-unicast | flowspec | unicast} | ipv6 {labeled-unicast | flowspec | unicast}
| | }
no address-family
```

Syntax Description

| | |
|-----------------------------|----------------------------------------------------------------------------------------------------------------|
| ipv4 unicast | Specifies IP Version 4 (IPv4) unicast address prefixes. |
| ipv4 labeled-unicast | Specifies IPv4 labeled-unicast address prefixes. This option is available in IPv4 neighbor configuration mode. |
| ipv6 unicast | Specifies IP Version 6 (IPv6) unicast address prefixes. |
| ipv6 labeled-unicast | Specifies IPv6 labeled-unicast address prefixes. This option is available in IPv6 neighbor configuration mode. |

Command Default

An address family must be explicitly configured in the router configuration mode for the address family to be active in BGP. Similarly, an address family must be configured under the neighbor for the BGP session to be established for that address family. An address family must be configured in router configuration mode before it can be configured under a neighbor.

Command Modes

Router configuration
 Neighbor configuration
 Neighbor group configuration
 Flowspec configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| 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.

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

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

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, , and labeled-unicast address families. IPv6 neighbor sessions support IPv6 unicast address families.

| Task ID | Task ID | Operations |
|---------|---------|----------------|
| | bgp | read, write |

Examples

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

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

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

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

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 | <i>seconds</i> 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 |
|----------------------|---------------------------------------------------------------------------------------|

| | |
|------------------------|--------------------------------------------|
| Command History | Release Modification |
| | Release 3.9.0 No modification. |
| | 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/RP0/CPU0:router (config) # router bgp 5
RP/0/RP0/CPU0:router (config-bgp) # neighbor 10.1.1.1
```

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

Related Commands

| Command | Description |
|---------------------------------------------|------------------------------------------------------------------------|
| neighbor-group, on page 141 | Creates a neighbor group and enters neighbor group configuration mode. |
| session-group, on page 189 | 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 | | |
|--------------------|-----------------------------|---------------------------------------------------------|
| | <i>af-group-name</i> | Address family group name. |
| | address-family | Enters address family configuration mode. |
| | ipv4 unicast | Specifies IP Version 4 (IPv4) unicast address prefixes. |
| | ipv4 labeled-unicast | Specifies IPv4 labeled unicast address prefixes. |
| | ipv6 unicast | Specifies IP Version 6 (IPv6) unicast address prefixes. |
| | ipv6 labeled-unicast | Specifies IPv6 labeled unicast address prefixes. |

Command Default No BGP address family group is configured.

Command Modes Router configuration

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | 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.

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.

| Task ID | Task ID | Operations |
|---------|---------|----------------|
| | bgp | read, write |

Examples

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

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

Related Commands

| Command | Description |
|---------------------------------------------|---------------------------------------------------------------------------------------|
| neighbor (BGP), on page 139 | Enters neighbor configuration mode for configuring BGP routing sessions. |
| neighbor-group, on page 141 | Creates a neighbor group and enters neighbor group configuration mode. |
| session-group, on page 189 | Creates a session group and enters session group configuration mode. |
| use, on page 368 | 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]
no aggregate-address
```

| Syntax Description | | |
|--------------------|----------------------------------------------|----------------------------------------------------------------------------------------------------------------|
| | <i>address</i> | Aggregate address. |
| | <i>/mask-length</i> | 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 <i>route-policy-name</i> | (Optional) Specifies the name of a route policy used to set the attributes of the aggregate route. |

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

Command Modes IPv4 address family configuration
IPv6 address family configuration

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | Release 3.9.0 | No modification. |
| | 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.

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

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

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

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

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

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

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

| Task ID | Task ID | Operations |
|---------|---------|----------------|
| | bgp | read, write |

Examples

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

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

| Related Commands | Command | Description |
|------------------|-------------------------------------------------|-----------------------------------------------------------------------------------|
| | network (BGP), on page 143 | Specifies the list of networks for the BGP routing process. |
| | route-policy (BGP), on page 178 | 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 receive of AiGP attribute is disabled for eBGP neighbors |
|------------------------|------------------------------------------------------------------|

| | |
|----------------------|---------------------------------------|
| Command Modes | IPv4 address family configuration |
| | IPv6 address family configuration |
| | Neighbor address family configuration |

| | | |
|------------------------|----------------|------------------------------|
| Command History | Release | Modification |
| | 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 |
| | bgp | read, write |

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

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

| | | |
|-------------------------|---------------------------------|-------------------------------------|
| 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 | | |
|---------------------|--|------------------------------------------------------------------------|
| <i>cost-comm-id</i> | | 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

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | 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.

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

| Task ID | Task ID | Operation |
|---------|---------|----------------|
| | bgp | read, write |

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

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

Related Commands

| Command | Description |
|----------------------------------|------------------------------------------------------------------------------------------|
| aigp, on page 14 | Enables sending and receiving of accumulated interior gateway protocol (AiGP) attribute. |

allocate-label

To allocate Multiprotocol Label Switching (MPLS) labels for specific 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 <i>route-policy-name</i> | Uses a route policy to select prefixes for label allocation. |

Command Default

No default behavior or values

Command Modes

IPv4 address family configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| 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.

Use the **allocate-label** command with a route policy to trigger BGP to allocate labels for all or a filtered set of global routes (as dictated by the route policy).



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.

Task ID

| Task ID | Operations |
|---------|----------------|
| bgp | read, write |

Examples

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

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

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/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router bgp 1
RP/0/RP0/CPU0:router(config-bgp)# vrf foo
RP/0/RP0/CPU0:router(config-bgp-vrf)# address-family ipv4 unicast
RP/0/RP0/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 | <i>as-occurrence-number</i> (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 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 | Operations |
| | bgp | read, write |

as-format

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

```
as-format asdot
no
```

Syntax Description

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

Command Default

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

Command Modes

XR Config

Command History

| Release | Modification |
|---------------|------------------------------|
| 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 | Operations |
|---------|----------------|
| bgp | read, write |

Examples

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

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


as-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

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | 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.

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

| Task ID | Task ID | Operation |
|---------|---------|----------------|
| | bgp | read, write |

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

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

attribute-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 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.

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

Task ID

| Task ID | Operation |
|---------|----------------|
| bgp | read, write |

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

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

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

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

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

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 <i>value</i> | Specifies the BFD session's multiplier value for the neighbor. |
|--------------------|--------------------------------------|----------------------------------------------------------------------|
| | minimum-interval <i>value</i> | 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 3.9.0 | No modification. |
| | 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 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) |

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

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 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 | Operations |
|----------------|----------------|-------------------|
| | bgp | read, write |

Examples

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

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

bgp attribute-download

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

bgp attribute-download
no bgp attribute-download

Syntax Description This command has no keywords or arguments.

Command Default BGP attribute download is not enabled.

Command Modes IPv4 unicast address family configuration

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | Release 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.

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

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

| Task ID | Task ID | Operations |
|---------|---------|----------------|
| | bgp | read, write |

Examples

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

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

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

RP/0/RP0/CPU0:router (config)# router bgp 50
```

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

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


bgp auto-policy-soft-reset disable

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

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

Syntax Description This command has no keywords or arguments.

Command Default Automatic soft reset of peers is enabled.

Command Modes Router configuration

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | 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.



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:

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

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

Command History

| Release | Modification |
|---------------|------------------------------|
| 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.

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

Task ID

| Task ID | Operations |
|---------|----------------|
| bgp | read, write |

Examples

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

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

Related Commands

| Command | Description |
|------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------|
| bgp bestpath compare-routerid , on page 32 | 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 34 | Allows the comparison of the Multi Exit Discriminator (MED) for paths from neighbors in different autonomous systems. |
| bgp bestpath med confed, on page 36 | Enables MED comparison among paths learned from confederation peers. |
| bgp bestpath med missing-as-worst, on page 38 | 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

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | 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.

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

| Task ID | Task ID | Operations |
|---------|---------|----------------|
| | bgp | read, write |

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

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

| Related Commands | Command | Description |
|------------------|---------------------------------------|--------------------------------------------|
| | show bgp, on page 193 | 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 | |
| Command History | Release | Modification |
| | 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.

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

| Task ID | Task ID | Operations |
|---------|---------|----------------|
| | bgp | read, write |

Examples

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

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

Related Commands

| Command | Description |
|---------------------------------------|--------------------------------------------|
| show bgp, on page 193 | 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 | |
| Command History | Release | Modification |
| | 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.

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

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

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

| Task ID | Task ID | Operations |
|----------------|----------------|-------------------|
| | bgp | read, write |

Examples

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

```
RP/0/RP0/CPU0:router(config)# router bgp 100  
RP/0/RP0/CPU0:router(config-bgp)# bgp bestpath med always
```

Related Commands

| Command | Description |
|---------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| bgp bestpath med confed, on page 36 | Enables MED comparison among paths learned from confederation peers. |
| bgp bestpath med missing-as-worst, on page 38 | 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 193 | 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 5.0.0 | This command was introduced. |
| Usage Guidelines | <p>To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.</p> <p>By default, the MED of the following paths is not compared with the MED of any other path:</p> <ul style="list-style-type: none"> • 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 <p>Use the bgp bestpath med confed command to affect how the following types of paths are treated in the BGP best-path algorithm:</p> <ul style="list-style-type: none"> • Paths containing only confederation segments • Paths containing confederation segments followed by an AS_SET <p>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.</p> <p>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:</p> <ul style="list-style-type: none"> • 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 <p>If the bgp bestpath med confed command is enabled, the software selects path 1 as the best path because it:</p> <ul style="list-style-type: none"> • Has a lower MED than path 2 | |

- Has a lower IGP metric than path 3

The MED is not compared with path 3 because it has an external autonomous system number (that is, an AS_SEQUENCE) in the path. If the **bgp bestpath med confed** command is not enabled, then MED is not compared between any of these paths. Consequently, the software selects path 2 as the best path because it has the lowest IGP metric.

| Task ID | Task ID | Operations |
|---------|---------|----------------|
| | bgp | read, write |

Examples

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

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

Related Commands

| Command | Description |
|---------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| bgp bestpath med always, on page 34 | Enables MED comparison among paths from neighbors in different autonomous systems. |
| bgp bestpath med missing-as-worst, on page 38 | 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 193 | 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

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | 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 | Operations |
|---------|---------|----------------|
| | bgp | read, write |

Examples

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

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

Related Commands

| Command | Description |
|-----------------------------------------------------|------------------------------------------------------------------------------------|
| bgp bestpath med always, on page 34 | Enables MED comparison among paths from neighbors in different autonomous systems. |
| bgp bestpath med confed, on page 36 | Enables MED comparison among paths learned from confederation peers. |
| show bgp, on page 193 | 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/RP0/CPU0:router#configure
RP/0/RP0/CPU0:router(config)#router bgp 50000
RP/0/RP0/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/RP0/CPU0:router#configure
RP/0/RP0/CPU0:router(config)#router bgp 50000
RP/0/RP0/CPU0:router(config-bgp)#address-family ipv4 unicast
RP/0/RP0/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/RP0/CPU0:router#configure
RP/0/RP0/CPU0:router(config)#router bgp 50000
RP/0/RP0/CPU0:router(config-bgp)#neighbor 1.1.1.1
RP/0/RP0/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/RP0/CPU0:router#configure
RP/0/RP0/CPU0:router(config)#router bgp 50000
RP/0/RP0/CPU0:router(config-bgp)#neighbor 1.1.1.1
RP/0/RP0/CPU0:router(config-bgp-nbr)#address-family ipv4 unicast
RP/0/RP0/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'.

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/RP0/CPU0:router#configure
RP/0/RP0/CPU0:router(config)#router bgp 50000
RP/0/RP0/CPU0:router(config-bgp)#bgp bestpath origin-as use validity
```

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/RP0/CPU0:router#configure
RP/0/RP0/CPU0:router(config)#router bgp 50000
RP/0/RP0/CPU0:router(config-bgp)#bgp bestpath aigp ignore
```

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/RP0/CPU0:router#configure
RP/0/RP0/CPU0:router(config)#router bgp 120
RP/0/RP0/CPU0:router(config-bgp)#bgp bestpath as-path multipath-relax
```

bgp client-to-client reflection disable

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

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

| Syntax Description | cluster-id <i>cluster-id</i> (Optional) Cluster ID for which intra-cluster route reflection is to be disabled; maximum of 4 bytes. Cluster ID can be entered either as an IP address or value. Range is 1 to 4294967295. | | | | |
|---------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------|--------------|---------------|------------------------------|
| Command Default | Client-to-client reflection is enabled. | | | | |
| Command Modes | Address family configuration | | | | |
| Command History | <table border="1"> <thead> <tr> <th>Release</th> <th>Modification</th> </tr> </thead> <tbody> <tr> <td>Release 5.0.0</td> <td>This command was introduced.</td> </tr> </tbody> </table> | Release | Modification | Release 5.0.0 | This command was introduced. |
| Release | Modification | | | | |
| 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.

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

Examples

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

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

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


Related Commands

| Command | Description |
|-----------------------------------------------------|-----------------------------------------------------------------------------------------------------|
| bgp cluster-id, on page 46 | Configures the cluster ID if the BGP cluster has more than one route reflector. |
| route-reflector-client, on page 180 | Configures the router as a BGP route reflector and configures the specified neighbor as its client. |
| show bgp, on page 193 | 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 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. |
|-------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

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

Related Commands

| Command | Description |
|-----------------------------------------------------|-----------------------------------------------------------------------------------------------------|
| cluster-id, on page 104 | Configures the cluster to which a neighbor belongs. |
| route-reflector-client, on page 180 | Configures the router as a BGP route reflector and configures the specified neighbor as its client. |
| show bgp, on page 193 | 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 | <p>as-number Autonomous system (AS) number that internally includes multiple autonomous systems.</p> <ul style="list-style-type: none"> • 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. |
|---------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

| | |
|------------------------|--------------------------------------------|
| Command Default | No confederation identifier is configured. |
|------------------------|--------------------------------------------|

| | |
|----------------------|----------------------|
| Command Modes | Router configuration |
|----------------------|----------------------|

| Command History | Release | Modification |
|------------------------|----------------|------------------------------|
| | Release 5.0.0 | This command was introduced. |

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

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

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

| Task ID | Task ID | Operations |
|----------------|----------------|-------------------|
| | bgp | read, write |

Examples

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

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

Related Commands

| Command | Description |
|-----------------------------------------------------|---------------------------------------------------------------------|
| bgp confederation peers, on page 50 | 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 | <p><i>as-number</i> Autonomous system (AS) numbers for Border Gateway Protocol (BGP) peers that belong to the confederation.</p> <ul style="list-style-type: none"> • 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. |
|---------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

| | |
|------------------------|----------------------------------------------------------------|
| Command Default | No BGP peers are identified as belonging to the confederation. |
|------------------------|----------------------------------------------------------------|

| | |
|----------------------|----------------------|
| Command Modes | Router configuration |
|----------------------|----------------------|

| | | |
|------------------------|----------------|------------------------------|
| Command History | Release | Modification |
| | Release 5.0.0 | This command was introduced. |

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

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 48](#) command specifies the confederation to which the autonomous systems belong.

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

| | | |
|----------------|----------------|-------------------|
| Task ID | Task ID | Operations |
| | bgp | read, write |

Examples

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

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

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

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

Related Commands

| Command | Description |
|----------------------------------------------------------|-------------------------------------------|
| bgp confederation identifier, on page 48 | 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 <i>half-life</i> value is allowed to default, the maximum suppress time defaults to 60 minutes. | |
| route-policy <i>route-policy-name</i> | (Optional) Specifies the route policy to use to set dampening parameters. | |

| Command Default | |
|--------------------------------------------------------------|--|
| Route dampening is disabled. | |
| <i>half-life</i> : 15 minutes | |
| <i>reuse</i> : 750 | |
| <i>suppress</i> : 2000 | |
| <i>max-suppress-time</i> : four times <i>half-life</i> value | |

| Command Modes | |
|-----------------------------------|--|
| IPv4 address family configuration | |
| IPv6 address family configuration | |

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | 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. | |

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

| Task ID | Task ID | Operations |
|---------|---------|----------------|
| | bgp | read, write |

Examples

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

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

Related Commands

| Command | Description |
|-------------------------------------------------------|--------------------------------------------------------------------------------|
| clear bgp dampening, on page 85 | Clears BGP route dampening information and unsuppresses the suppressed routes. |
| clear bgp flap-statistics, on page 88 | Clears BGP flap statistics. |
| route-policy (RPL) | Defines a route policy and enters route-policy configuration mode. |
| show bgp dampened-paths, on page 238 | Displays BGP dampened routes. |
| show bgp flap-statistics, on page 242 | Displays BGP flap statistics. |
| show bgp neighbors, on page 256 | 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

Command History

| Release | Modification |
|---------------|------------------------------|
| 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.

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

Task ID

| Task ID | Operations |
|---------|----------------|
| bgp | read, write |

Examples

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

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

bgp enforce-first-as disable

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

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

| | | |
|---------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------|
| Syntax Description | This command has no keywords or arguments. | |
| Command Default | By default, the software requires the first autonomous system (in the AS path) of a route received from an eBGP peer to be the same as the remote autonomous system configured. | |
| Command Modes | Router configuration | |
| Command History | Release | Modification |
| | Release 3.9.0 | No modification. |
| | Release 5.0.0 | This command was introduced. |
| Usage Guidelines | <p>To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.</p> <p>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.</p> | |
| Task ID | Task ID | Operations |
| | bgp | read, write |
| Examples | <p>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:</p> <pre>RP/0/RP0/CPU0:router(config)# router bgp 100 RP/0/RP0/CPU0:router(config-bgp)# bgp enforce-first-as disable</pre> | |

| Related Commands | Command | Description |
|------------------|-------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| | enforce-first-as, on page 118 | Disables the software to enforce the first autonomous system in 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, in neighbor configuration mode, neighbor group configuration mode, and session group configuration mode. |
| | enforce-first-as-disable, on page 120 | Disables the software to enforce the first autonomous system in 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, in neighbor configuration mode, neighbor group configuration mode, and session group configuration mode. |
| | show bgp, on page 193 | 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 fallover. | | | | |
|---------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------|--------------|---------------|------------------------------|
| 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 | | | | |
| Command History | <table border="1"> <thead> <tr> <th>Release</th> <th>Modification</th> </tr> </thead> <tbody> <tr> <td>Release 5.0.0</td> <td>This command was introduced.</td> </tr> </tbody> </table> | Release | Modification | Release 5.0.0 | This command was introduced. |
| Release | Modification | | | | |
| Release 5.0.0 | This command was introduced. | | | | |
| Usage Guidelines | <p>To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.</p> <p>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.</p> | | | | |
| Task ID | <table border="1"> <thead> <tr> <th>Task ID</th> <th>Operations</th> </tr> </thead> <tbody> <tr> <td>bgp</td> <td>read, write</td> </tr> </tbody> </table> | Task ID | Operations | bgp | read, write |
| Task ID | Operations | | | | |
| bgp | read, write | | | | |

Examples

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

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

bgp graceful-restart

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

bgp graceful-restart
no bgp graceful-restart

Syntax Description This command has no keywords or arguments.

Command Default Graceful restart support is not enabled.

Command Modes Router configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 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.

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.

Task ID

| Task ID | Operations |
|---------|----------------|
| bgp | read, write |

Examples

The following example shows how to enable graceful restart:

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

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

Related Commands

| Command | Description |
|-----------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------|
| bgp graceful-restart graceful-reset, on page 60 | Enables a graceful reset if configuration changes force a peer reset. |
| bgp graceful-restart purge-time, on page 61 | Defines the maximum time before stale routes are purged. |
| bgp graceful-restart restart-time, on page 62 | Defines the maximum time advertised to neighbors |
| bgp graceful-restart stalepath-time, on page 63 | 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 193 | Displays entries in the BGP routing table. |
| show bgp neighbors, on page 256 | Displays information about BGP connections to neighbors. |
| show bgp process, on page 290 | 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 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.

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

| Task ID | Task ID | Operations |
|----------------|----------------|-------------------|
| | bgp | read, write |

Examples

The following example shows how to enable graceful reset:

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

| Related Commands | Command | Description |
|-------------------------|--------------------------------------------------|----------------------------------------------------------|
| | bgp graceful-restart, on page 58 | Enables a graceful restart. |
| | show bgp, on page 193 | Displays entries in the BGP routing table. |
| | show bgp neighbors, on page 256 | Displays information about BGP connections to neighbors. |
| | show bgp process, on page 290 | 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 | <i>seconds</i> Maximum time before stale routes are purged. Time in seconds. Range is 0 to 6000. |
|---------------------------|--------------------------------------------------------------------------------------------------|

| | |
|------------------------|----------------------|
| Command Default | <i>seconds</i> : 600 |
|------------------------|----------------------|

| | |
|----------------------|----------------------|
| Command Modes | Router configuration |
|----------------------|----------------------|

| Command History | Release | Modification |
|------------------------|----------------|------------------------------|
| | 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.

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

| Task ID | Task ID | Operations |
|----------------|----------------|-------------------|
| | bgp | read, write |

Examples

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

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

| Related Commands | Command | Description |
|-------------------------|--------------------------------------------------|----------------------------------------------------------|
| | bgp graceful-restart, on page 58 | Enables a graceful restart. |
| | show bgp, on page 193 | Displays entries in the BGP routing table. |
| | show bgp neighbors, on page 256 | Displays information about BGP connections to neighbors. |
| | show bgp process, on page 290 | 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 | <i>seconds</i> Maximum time advertised to neighbors. Time in seconds. Range is 1 to 4095. |
|---------------------------|-------------------------------------------------------------------------------------------|

| | |
|------------------------|----------------------|
| Command Default | <i>seconds</i> : 120 |
|------------------------|----------------------|

| | |
|----------------------|----------------------|
| Command Modes | Router configuration |
|----------------------|----------------------|

| Command History | Release | Modification |
|------------------------|----------------|------------------------------|
| | 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. |
|-------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

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

| Task ID | Task ID | Operations |
|----------------|----------------|-------------------|
| | bgp | read, write |

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

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

| Related Commands | Command | Description |
|-------------------------|--------------------------------------------------|----------------------------------------------------------|
| | bgp graceful-restart, on page 58 | Enables a graceful restart. |
| | show bgp, on page 193 | Displays entries in the BGP routing table. |
| | show bgp neighbors, on page 256 | Displays information about BGP connections to neighbors. |
| | show bgp process, on page 290 | 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 | <i>seconds</i> Maximum wait time. Time in seconds. Range is 1 to 4095. |
|---------------------------|------------------------------------------------------------------------|

| | |
|------------------------|----------------------|
| Command Default | <i>seconds</i> : 360 |
|------------------------|----------------------|

| | |
|----------------------|----------------------|
| Command Modes | Router configuration |
|----------------------|----------------------|

| Command History | Release | Modification |
|------------------------|----------------|------------------------------|
| | 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.

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

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

| Task ID | Task ID | Operations |
|----------------|----------------|-------------------|
| | bgp | read, write |

Examples

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

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

| Related Commands | Command | Description |
|-------------------------|--------------------------------------------------|----------------------------------------------------------|
| | bgp graceful-restart, on page 58 | Enables a graceful restart. |
| | show bgp, on page 193 | Displays entries in the BGP routing table. |
| | show bgp neighbors, on page 256 | Displays information about BGP connections to neighbors. |

| Command | Description |
|-----------------------------------------------|-----------------------------------|
| show bgp process, on page 290 | Displays BGP process information. |

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 | <i>seconds</i> | Specifies batch label processing delay in seconds. Range is 0 to 10 seconds. |
| | <i>milliseconds</i> | 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 | |
| Command History | Release | Modification |
| | 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 |
| | bgp | read, write |

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

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

bgp 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

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | 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.

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

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

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

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

| Task ID | Task ID | Operations |
|---------|---------|----------------|
| | bgp | read, write |

Examples

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

```
RP/0/RP0/CPU0:router(config)# router bgp 65530
```

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

Related Commands

| Command | Description |
|-------------------------------------------------|----------------------------------------------------------------------|
| show bgp neighbors, on page 256 | Displays information about the TCP and BGP connections to neighbors. |

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 | <i>limit</i> Maximum number of neighbors. Range is 1 to 15000. | |
| Command Default | Default limit is 4000 | |
| Command Modes | Router configuration | |
| Command History | Release | Modification |
| | 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.

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

| | | |
|----------------|----------------|-------------------|
| Task ID | Task ID | Operations |
| | bgp | write |

Examples

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

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


bgp 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
```

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

Command History

| Release | Modification |
|---------------|------------------------------|
| 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.

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



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

Task ID

| Task ID | Operations |
|---------|----------------|
| bgp | read, write |

Examples

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

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

Related Commands

| Command | Description |
|-----------------------------------------|---------------------------|
| clear bgp, on page 82 * | 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 | <i>ip-address</i> 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 | | | | |
| Command History | <table border="1"> <thead> <tr> <th>Release</th> <th>Modification</th> </tr> </thead> <tbody> <tr> <td>Release 5.0.0</td> <td>This command was introduced.</td> </tr> </tbody> </table> | Release | Modification | Release 5.0.0 | This command was introduced. |
| Release | Modification | | | | |
| Release 5.0.0 | This command was introduced. | | | | |
| Usage Guidelines | <p>To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.</p> <p>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.</p> <p>For more details on router IDs, see the <i>Routing Configuration Guide for Cisco NCS 6000 Series Routers</i></p> | | | | |
| Task ID | <table border="1"> <thead> <tr> <th>Task ID</th> <th>Operations</th> </tr> </thead> <tbody> <tr> <td>bgp</td> <td>read, write</td> </tr> </tbody> </table> | Task ID | Operations | bgp | read, write |
| Task ID | Operations | | | | |
| bgp | read, write | | | | |

Examples

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

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

Related Commands

| Command | Description |
|---------------------------------------|--------------------------------------------|
| show bgp, on page 193 | 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 | <i>seconds</i> 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 |
|----------------------|------------------------------------------------------------------------------------------------|

| Command History | Release | Modification |
|------------------------|----------------|------------------------------|
| | 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. |
|-------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

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

| Task ID | Task ID | Operations |
|----------------|----------------|-------------------|
| | bgp | read, write |

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

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

| Related Commands | Command | Description |
|-------------------------|---------------------------------------|--------------------------------------------|
| | show bgp, on page 193 | Displays entries in the BGP routing table. |

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 | <i>seconds</i> 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 | <table border="1"> <thead> <tr> <th>Release</th> <th>Modification</th> </tr> </thead> <tbody> <tr> <td>Release 5.0.0</td> <td>This command was introduced.</td> </tr> </tbody> </table> | Release | Modification | Release 5.0.0 | This command was introduced. |
| Release | Modification | | | | |
| 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.

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.

| Task ID | Task ID | Operations |
|---------|---------|----------------|
| | bgp | read, write |

Examples

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

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

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 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.

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.

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

| | | |
|----------------|----------------|------------------|
| Task ID | Task ID | Operation |
| | bgp | read, write |

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

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

capability additional-paths send

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

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

| | |
|---------------------------|--------------------------------------------------------------------|
| Syntax Description | disable Disables advertise additional paths send capability |
|---------------------------|--------------------------------------------------------------------|

| | |
|------------------------|-------------------------|
| Command Default | Capability is disabled. |
|------------------------|-------------------------|

| | |
|----------------------|---------------------------------------------------------------------------------------|
| Command Modes | Neighbor configuration Neighbor group configuration Session group configuration |
|----------------------|---------------------------------------------------------------------------------------|

| | | |
|------------------------|----------------|------------------------------|
| Command History | Release | Modification |
| | Release 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. |
|-------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

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.

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

| | | |
|----------------|----------------|------------------|
| Task ID | Task ID | Operation |
| | bgp | read, write |

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

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


capability orf prefix

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

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

Syntax Description

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

Command Default

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

Command Modes

IPv4 address family group configuration
 IPv6 address family group configuration
 IPv4 neighbor address family configuration
 IPv4 neighbor group address family configuration
 IPv6 neighbor group address family configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| 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.

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



Note

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

If the neighbor sends a prefix list and the speaker accepts it, the speaker applies the received prefix list, plus any locally configured outbound filters, to limit its outbound routing updates to the neighbor. Increased filtering

prevents unwanted routing updates between neighbors and reduces resource requirements for routing update generation and processing.

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



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

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

Task ID

| Task ID | Operations |
|---------|----------------|
| bgp | read, write |

Examples

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

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

Related Commands

| Command | Description |
|-------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------|
| af-group, on page 10 | Creates an address family group for BGP neighbors and enters address family group configuration mode. |
| neighbor-group, on page 141 | Creates a neighbor group and enters neighbor group configuration mode. |
| show bgp neighbors, on page 256 | 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-nextthop-encoding} [**inheritance-disable**]

| Syntax Description | 4-byte-as | extended-nextthop-encoding | inheritance-disable |
|--------------------|--------------------------------------------------------------------|------------------------------------------------------------|---------------------------------------------------------------|
| | 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 5.0.0 | This command was introduced. |

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



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

| Task ID | Task ID | Operations |
|---------|---------|----------------|
| | bgp | read, write |

Examples

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

```
RP/0/RP0/CPU0:router# show bgp nei 10.3.3.3 conf
neighbor 10.3.3.3
  remote-as 65000                [n:internal]
  description PE3                []
```

```
update-source Loopback0 [n:internal]
address-family ipv4 unicast [n:internal]
```

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

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

```
For Address Family: IPv4 Unicast
```

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

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

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

```
RP/0/RP0/CPU0:router(config-bgp)#neighbor 10.3.3.3
```

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

```
RP/0/RP0/CPU0:router(config-bgp-nbr)#commit
```

```
RP/0/RP0/CPU0:router(config-bgp-nbr)#end
```

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

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

```
For Address Family: IPv4 Unicast
```

```
BGP neighbor version 1
Update group: 0.2
Route refresh request: received 0, sent 0
0 accepted prefixes, 0 are bestpaths
Prefix advertised 0, suppressed 0, withdrawn 0, maximum limit 524288
Threshold for warning message 75%
```

An EoR was received during read-only mode

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

With the **inheritance-disable** keyword:

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

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

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

clear bgp

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

```
clear bgp [{ipv4 {unicast | labeled-unicast | all} | ipv6 {unicast | all | labeled-unicast} | all }]
```

| Syntax Description | |
|------------------------|-----------------------------------------------------------------------------|
| ipv4 | (Optional) Specifies IP Version 4 address prefixes. |
| unicast | (Optional) Specifies unicast address prefixes. |
| labeled-unicast | (Optional) Specifies labeled unicast address prefixes. |
| ipv6 | (Optional) Specifies IP Version 6 address prefixes. |
| all | (Optional) For address family, specifies prefixes for all address families. |

Command Default No default behavior or values

Command Modes XR EXEC

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | 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.

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

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

| Task ID | Task ID | Operations |
|---------|---------|------------|
| | bgp | execute |

Examples

The following example shows how to hard reset neighbor 10.0.0.1:

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

Related Commands

| Command | Description |
|-------------------------------------------------------|----------------------------------------------------------------------|
| clear bgp self-originated, on page 98 | Clears self-originated routes. |
| clear bgp soft, on page 101 | Soft resets a group of BGP neighbors. |
| show bgp, on page 193 | Displays entries in the BGP routing table. |
| show bgp neighbors, on page 256 | Displays information about the TCP and BGP connections to neighbors. |

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
```


clear bgp dampening

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

clear bgp dampening

| Syntax Description | | |
|------------------------|--|----------------------------------------------------------------------------------|
| ipv4 | | Specifies IP Version 4 address prefixes. |
| unicast | | Specifies unicast address prefixes. |
| labeled-unicast | | Specifies labeled-unicast address prefixes. |
| ipv6 | | Specifies IP Version 6 address prefixes. |
| all | | For address family, specifies prefixes for all address families. |
| <i>ip-address</i> | | (Optional) IP address of the network about which to clear dampening information. |
| <i>/mask-length</i> | | (Optional) Network mask applied to the IP address. |

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

Command Modes XR EXEC

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | 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.

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

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

| Task ID | Task ID | Operations |
|---------|---------|------------|
| | bgp | execute |

Examples

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

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

Related Commands

| Command | Description |
|------------------------------------------------------|-----------------------------------------------------------------------------|
| bgp dampening, on page 52 | Enables BGP route dampening or changes various BGP route dampening factors. |
| show bgp dampened-paths, on page 238 | Displays BGP dampened routes. |

clear bgp external

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

clear bgp external

| Syntax Description | | |
|------------------------|------------|---------------------------------------------------------------------------------------------------------------------------------------|
| ipv4 | (Optional) | Specifies IP Version 4 address prefixes. |
| unicast | (Optional) | Specifies unicast address prefixes. |
| labeled-unicast | (Optional) | Specifies labeled unicast address prefixes. |
| ipv6 | (Optional) | Specifies IP Version 6 address prefixes. |
| all | (Optional) | For address family, specifies prefixes for all address families. |
| graceful | (Optional) | Clears all external peers with a hard reset and a graceful restart. This option is available when an address family is not specified. |

Command Default No default behavior or value

Command Modes XR EXEC

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | 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 | Operations |
|---------|---------|------------|
| | bgp | execute |

Examples

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

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

clear bgp flap-statistics

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

clear bgp flap-statistics

| Syntax Description | | |
|----------------------------------------------|--|--------------------------------------------------------------------------------------------------|
| ipv4 | | Specifies IP Version 4 address prefixes. |
| unicast | | Specifies unicast address prefixes. |
| 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. |
| regexp <i>regexp</i> | | (Optional) Clears flap statistics for routes whose AS paths match the regular expression. |
| route-policy <i>route-policy-name</i> | | (Optional) Clears flap statistics for the specific route policy. |
| <i>network</i> | | (Optional) Network for which flap counts are to be cleared. |
| <i>/mask-length</i> | | (Optional) Network mask of the network for which flap counts are to be cleared. |
| <i>ip-address</i> | | (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 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 | Operations |
|---------|---------|------------|
| | bgp | execute |

Examples

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

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

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

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

clear bgp nexthop performance-statistics

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

clear bgp nexthop performance-statistics

| Syntax Description | Option | Description |
|--------------------|------------------------|--------------------------------------------------------------------------|
| | ipv4 | Specifies IP Version 4 address prefixes. |
| | unicast | Specifies unicast 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. |

Command Default No default behavior or values

Command Modes XR EXEC

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | 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.

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 |

Related Commands

| Command | Description |
|------------------------------------------------|------------------------------------------------------------|
| show bgp nexthops, on page 275 | 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 XR EXEC mode.

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

| Syntax Description | | |
|--------------------|------------------------|--------------------------------------------------------------------------|
| | ipv4 | Specifies IP Version 4 address prefixes. |
| | unicast | Specifies unicast address prefixes. |
| | 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. |
| | <i>nexthop-address</i> | Address of the next hop. |

Command Default No default behavior or values

Command Modes XR EXEC

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | 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.

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

| Task ID | Task ID | Operations |
|---------|---------|------------|
| | bgp | execute |

Examples

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

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


Related Commands

| Command | Description |
|------------------------------------------------|------------------------------------------------------------|
| show bgp nexthops, on page 275 | 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 XR EXEC mode.

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

| Syntax Description | |
|--------------------|--------------------------------------------|
| * | Specifies all BGP neighbors. |
| <i>ip-address</i> | IP address of a specific network neighbor. |

Command Default No default behavior or values

Command Modes XR EXEC

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | 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 | Operations |
|---------|---------|------------|
| | bgp | execute |

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

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

| Related Commands | Command | Description |
|------------------|-------------------------------------------------|----------------------------------------------------------|
| | show bgp neighbors, on page 256 | 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 performance-statistics
```

Syntax Description This command has no keywords or arguments.

Command Default No default behavior or values

Command Modes XR EXEC

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | 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 | Operations |
|---------|---------|------------|
| | bgp | execute |

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

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

clear bgp rpki

To clear resource public key infrastructure (RPKI) cache-server (rpki-server) 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 XR EXEC mode.

clear bgp self-originated

| Syntax Description | Option | Description |
|--------------------|------------------------|--------------------------------------------------------------------------|
| | ipv4 | Specifies IP Version 4 address prefixes. |
| | unicast | Specifies unicast 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. |

Command Default No default behavior or values

Command Modes XR EXEC

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | 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.

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

| Task ID | Task ID | Operations |
|---------|---------|------------|
| | bgp | execute |

Examples

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

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

| Related Commands | Command | Description |
|------------------|-----------------------------------------------|----------------------------------------------------|
| | aggregate-address, on page 12 | Creates an aggregate entry in a BGP routing table. |

| Command | Description |
|-------------------------------------------------|---------------------------------------------------------------------------------------------------------|
| network (BGP), on page 143 | Specifies a local network that the BGP routing process should originate and advertise to its neighbors. |
| redistribute (BGP), on page 169 | 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 XR EXEC mode.

clear bgp shutdown

| | |
|------------------------|--------------------------------------------------------------------------|
| ipv4 | Specifies IP Version 4 address prefixes. |
| unicast | Specifies unicast 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. |

Command Default No default behavior or values

Command Modes XR EXEC

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | 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 | Operations |
|---------|---------|------------|
| | bgp | execute |

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

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

| Related Commands | Command | Description |
|------------------|-------------------------------------------------|----------------------------------------------------------------------|
| | show bgp, on page 193 | Displays entries in the BGP routing table. |
| | show bgp neighbors, on page 256 | 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 XR EXEC mode.

```
clear bgp { * 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. |
| 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. |
| <i>nexthop-address</i> | | Address of the next hop. |
| * | | Soft resets all BGP neighbors. |
| <i>ip-address</i> | | IP address of the neighbor to be reset. |
| as <i>as-number</i> | | Autonomous system (AS) number for all neighbors to be reset. Range for 2-byte numbers is 1 to 65535. Range for 4-byte numbers is 1.0 to 65535.65535. |
| external | | Specifies clearing of all external peers. |
| in | | (Optional) Triggers an inbound soft reset. If the in or out keyword is not specified, both inbound and outbound soft resets are triggered. |
| prefix-filter | | (Optional) Specifies to send a new Outbound Route Filter (ORF) to the neighbor. Neighbor installs the new ORF and resends its routes. |
| out | | (Optional) Triggers an outbound soft reset. If the in or out keyword is not specified, both inbound and outbound soft resets are triggered. |

Command Default No default behavior or value

Command Modes XR EXEC

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | 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.

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 29 | Disables an automatic soft reset of BGP peers when the configured inbound route policy is modified. |
| clear bgp, on page 82 | Resets a group of BGP neighbors. |
| clear bgp self-originated, on page 98 | Clears self-originated routes. |
| show bgp, on page 193 | Displays entries in the BGP routing table. |

| Command | Description |
|-----------------------------------------------------------|----------------------------------------------------------------------|
| show bgp neighbors, on page 256 | Displays information about the TCP and BGP connections to neighbors. |
| soft-reconfiguration inbound, on page 345 | Configures the software to store updates received from a neighbor. |

cluster-id

To configure the cluster for a neighbor, use the **cluster-id** command in an appropriate configuration mode. To remove the cluster, use the **no** form of this command.

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

Syntax Description

cluster-id Cluster ID of the router acting as a route reflector; maximum of four 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

Neighbor configuration
 Neighbor group configuration
 Session group configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| 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.

A single route reflector can support multiple clusters. A neighbor can be associated with one cluster only. And the corresponding cluster ID is configured in neighbor configuration mode. If the cluster ID is not configured for a neighbor and the neighbor is a route reflector client, then the neighbor is assigned to the default cluster.

A neighbor will be considered to be a route reflector client only if it is configured as a route reflector client in the appropriate address-family configuration mode.

Configuring the cluster ID using the **cluster-id** command for a neighbor group or session group under the neighbor group configuration mode or the session group configuration mode 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.

To increase redundancy and avoid a single point of failure in the network, the clusters might be connected to more than one route reflector. In this case, the neighbor to cluster-id mapping at all the route reflectors must be the same so that a route reflector can recognize updates from route reflectors that are connected to the same clusters.

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 three clusters. Neighbor 192.168.70.25 is assigned to the default cluster with cluster ID 1.

```
RP/0/RP0/CPU0:router(config)# router bgp 65534
RP/0/RP0/CPU0:router(config-bgp)# bgp cluster-id 1
RP/0/RP0/CPU0:router(config-bgp)# neighbor 192.168.70.24
RP/0/RP0/CPU0:router(config-bgp-nbr)# remote-as 65534
RP/0/RP0/CPU0:router(config-bgp-nbr)# cluster-id 2
RP/0/RP0/CPU0:router(config-bgp-nbr)# address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-bgp-nbr-af)# route-reflector-client

RP/0/RP0/CPU0:router(config-bgp)# neighbor 192.168.70.25
RP/0/RP0/CPU0:router(config-bgp-nbr)# remote-as 65534
RP/0/RP0/CPU0:router(config-bgp-nbr)# address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-bgp-nbr-af)# route-reflector-client

RP/0/RP0/CPU0:router(config-bgp)# neighbor 192.168.70.26
RP/0/RP0/CPU0:router(config-bgp-nbr)# remote-as 65534
RP/0/RP0/CPU0:router(config-bgp-nbr)# cluster-id 3
RP/0/RP0/CPU0:router(config-bgp-nbr)# address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-bgp-nbr-af)# route-reflector-client
```

Related Commands

| Command | Description |
|-----------------------------------------------------|-----------------------------------------------------------------------------------------------------|
| route-reflector-client, on page 180 | Configures the router as a BGP route reflector and configures the specified neighbor as its client. |
| show bgp, on page 193 | Displays entries in the BGP routing table. |

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

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | 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.

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

| Task ID | Task ID | Operations |
|---------|---------|----------------|
| | bgp | read, write |

Examples

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

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

Related Commands

| Command | Description |
|-------------------------------------------------|------------------------------------------------------|
| redistribute (BGP), on page 169 | Redistributes routes from another protocol into BGP. |

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 | <i>value</i> 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 |
|----------------------|----------------------|

| Command History | Release | Modification |
|------------------------|----------------|------------------------------|
| | 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. |
|-------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

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



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

| Task ID | Task ID | Operations |
|----------------|----------------|-------------------|
| | bgp | read, write |

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

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


default-originate

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

```
default-originate [{inheritance-disable | route-policy route-policy-name}]
no default-originate [{inheritance-disable | route-policy route-policy-name}]
```

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

Command Default The default route is not advertised to BGP neighbors.

Command Modes

- IPv4 neighbor address family configuration
- IPv6 neighbor address family configuration
- IPv4 neighbor group address family configuration
- IPv6 neighbor group address family configuration
- IPv4 address family group configuration
- IPv6 address family group configuration

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | 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.

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

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

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

Related Commands

| Command | Description |
|------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------|
| default-information originate (BGP), on page 107 | Allows the default route to be redistributed into BGP from another routing protocol. |
| af-group, on page 10 | Creates an address family group for BGP neighbors and enters address family group configuration mode. |
| neighbor-group, on page 141 | Creates a neighbor group and enters neighbor group configuration mode. |

description (BGP)

To annotate a neighbor, neighbor group, 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

Command History

| Release | Modification |
|---------------|------------------------------|
| 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.

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

Task ID

| Task ID | Operations |
|---------|----------------|
| bgp | read, write |

Examples

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

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

distance bgp

To allow the use of external, internal, and local administrative distances that could be used to prefer one class of routes over another, use the **distance bgp** command in an appropriate configuration mode. To disable the use of administrative distances, use the **no** 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

Command History

| Release | Modification |
|---------------|------------------------------|
| 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.

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



Note

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

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

| Task ID | Task ID | Operations |
|---------|---------|----------------|
| | bgp | read, write |

Examples

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

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

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. |

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 | <i>value</i> 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 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. |
|-------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

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

[Table 2: dscp Default Settings, on page 114](#) 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/RP0/CPU0:router(config)# router bgp 5
RP/0/RP0/CPU0:router(config-bgp)#neighbor 10.1.1.1
RP/0/RP0/CPU0:router(config-bgp-nbr)#remote-as 100
RP/0/RP0/CPU0:router(config-bgp-nbr)# dscp af32
```

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

Neighbor group configuration

Session group configuration

Command History

Release

Modification

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.

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

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

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

Task ID

| Task ID | Operations |
|---------|------------|
|---------|------------|

| | |
|-----|----------------|
| bgp | read, write |
|-----|----------------|

Examples

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

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

Related Commands

| Command | Description |
|---------------------------------------------|------------------------------------------------------------------------|
| neighbor-group, on page 141 | Creates a neighbor group and enters neighbor group configuration mode. |
| session-group, on page 189 | Creates a session group and enters session group configuration mode. |

enforce-first-as

To enable the software to enforce the first autonomous system in 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 **enforce-first-as** command in an appropriate configuration mode. To disable enforcing 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, use the **no** form of this command.

enforce-first-as
no enforce-first-as

| | |
|---------------------------|--------------------------------------------|
| Syntax Description | This command has no arguments or keywords. |
|---------------------------|--------------------------------------------|

| | |
|------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 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 | Neighbor configuration Neighbor group configuration Session group configuration |
|----------------------|---------------------------------------------------------------------------------------|

| Command History | Release | Modification |
|------------------------|----------------|------------------------------|
| | 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. |
|-------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

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 under the neighbor, neighbor group or session group.

At any given time, either the **enforce-first-as** command or the [enforce-first-as-disable, on page 120](#) command can be configured under a given neighbor, neighbor group or session group. Configuring one command overwrites the other command.

| Task ID | Task ID | Operations |
|----------------|----------------|-------------------|
| | bgp | read, write |

| | |
|-----------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Examples | The following example shows a configuration in which incoming updates from eBGP neighbors are checked to ensure the first AS number in the AS path is the same as the configured AS number for the neighbor: |
|-----------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

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

```
RP/0/RP0/CPU0:router(config-bgp)# neighbor 10.2.3.4  
RP/0/RP0/CPU0:router(config-bgp-nbr)# enforce-first-as
```

| Related Commands | Command | Description |
|------------------|----------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| | bgp enforce-first-as disable, on page 55 | Disables the software to enforce the first autonomous system in 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, in router configuration mode mode. |
| | enforce-first-as-disable, on page 120 | Disables the software to enforce the first autonomous system in 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, in neighbor configuration mode, neighbor group configuration mode, and session group configuration mode. |

enforce-first-as-disable

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

enforce-first-as-disable

no enforce-first-as-disable

| | | |
|---------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------|
| Syntax Description | This command has no arguments or keywords. | |
| 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 | Neighbor configuration Neighbor group configuration Session group configuration | |
| Command History | Release | Modification |
| | Release 5.0.0 | This command was introduced. |
| Usage Guidelines | <p>To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.</p> <p>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 under the neighbor, neighbor-group or session-group.</p> <p>At any given time, either the enforce-first-as-disable command or the enforce-first-as, on page 118 command can be configured under a given neighbor, neighbor group or session group. Configuring one command overwrites the other command.</p> | |
| Task ID | Task ID | Operations |
| | bgp | read, write |
| Examples | <p>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:</p> <pre>RP/0/RP0/CPU0:router(config)# router bgp 100</pre> | |

```
RP/0/RP0/CPU0:router(config-bgp)# neighbor 10.2.3.4
RP/0/RP0/CPU0:router(config-bgp-nbr)# enforce-first-as-disable
```

Related Commands

| Command | Description |
|----------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| bgp enforce-first-as disable, on page 55 | Disables the software to enforce the first autonomous system in 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, in router configuration mode. |
| enforce-first-as, on page 118 | Enables the software to enforce the first autonomous system in 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, under neighbor configuration mode, neighbor group configuration mode, and session group configuration mode. |

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 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.

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

If the **ibgp policy out enforce-modifications** command is configured under router BGP configuration, then all the changes made by the outbound policy for an iBGP peer will be present in an update message sent to the peer.

Task ID

| Task ID | Operations |
|---------|----------------|
| bgp | read, write |

Examples

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

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router bgp 6500
RP/0/RP0/CPU0:router(config-bgp)# ibgp policy out enforce-modifications
```

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 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 | Operations |
| | bgp | read, write |


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

```
RP/0/RP0/CPU0:router(config)# router bgp 100
RP/0/RP0/CPU0:router(config-bgp)# neighbor 10.2.3.4
RP/0/RP0/CPU0:router(config-bgp-nbr)# ignore-connected-check
```

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 | <i>name</i> 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 | <table border="1"> <thead> <tr> <th>Release</th> <th>Modification</th> </tr> </thead> <tbody> <tr> <td>Release 5.0.0</td> <td>This command was introduced.</td> </tr> </tbody> </table> | Release | Modification | Release 5.0.0 | This command was introduced. |
| Release | Modification | | | | |
| Release 5.0.0 | This command was introduced. | | | | |
| Usage Guidelines | <p>To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.</p> <p>Specify a key chain to enable key chain authentication between two BGP peers. Use the keychain command to implement hitless key rollover for authentication.</p> <p>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.</p> | | | | |
|  Note | BGP only supports HMAC-MD5 and HMAC-SHA1-12 cryptographic algorithms. | | | | |
| Task ID | <table border="1"> <thead> <tr> <th>Task ID</th> <th>Operations</th> </tr> </thead> <tbody> <tr> <td>bgp</td> <td>read, write</td> </tr> </tbody> </table> | Task ID | Operations | bgp | read, write |
| Task ID | Operations | | | | |
| bgp | read, write | | | | |
| Examples | The following example shows how to configure neighbor 172.20.1.1 to use the key chain authentication configured in the keychain_A key chain: | | | | |


```
RP/0/RP0/CPU0:router(config)# router bgp 140  
RP/0/RP0/CPU0:router(config-bgp)# neighbor 172.20.1.1  
RP/0/RP0/CPU0:router(config-bgp-nbr)# remote-as 1  
RP/0/RP0/CPU0:router(config-bgp-nbr)# keychain keychain_A
```

Related Commands

| Command | Description |
|-----------------------------------------------|-----------------------------------------------------------------------------------------------------------|
| keychain-disable, on page 126 | 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 128](#) 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 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 you specify a key chain on a neighbor group or session group, all users of the group inherit the key chain. Specifying a different **keychain** command specifically on a neighbor that uses the group overrides the inherited value. Specifying **keychain-disable** on a neighbor that uses the group disables key chain authentication for the neighbor.

| Task ID | Task ID | Operations |
|---------|---------|----------------|
| | bgp | read, write |

Examples

The following example shows how to disable key chain authentication for neighbor 172.20.1.1, preventing it from inheriting the key chain keychain_A from session group group1:

```
RP/0/RP0/CPU0:router (config) # router bgp 140
```

```
RP/0/RP0/CPU0:router(config-bgp)# session-group group1
RP/0/RP0/CPU0:router(config-bgp-sngrp)# keychain keychain_A
RP/0/RP0/CPU0:router(config-bgp-sngrp)# exit
RP/0/RP0/CPU0:router(config-bgp)#neighbor 172.20.1.1
RP/0/RP0/CPU0:router(config-bgp-nbr)# remote-as 2
RP/0/RP0/CPU0:router(config-bgp-nbr)#use session-group group1
RP/0/RP0/CPU0:router(config-bgp-nbr)# keychain-disable
```

Related Commands

| Command | Description |
|---------------------------------------|---------------------------------------------------------------------------------|
| keychain, on page 124 | 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 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 you specify a key chain on a neighbor group or session group, all users of the group inherit the key chain. Specifying a different **keychain** command specifically on a neighbor that uses the group overrides the inherited value. Specifying **keychain inheritance-disable** on a neighbor that uses the group disables key chain authentication for the neighbor.

| Task ID | Task ID | Operations |
|----------------|----------------|-------------------|
| | bgp | read, write |

Examples

The following example shows how to disable key chain authentication for neighbor 172.20.1.1, preventing it from inheriting the key chain keychain_A from session group group1:

```
RP/0/RP0/CPU0:router (config) #router bgp 140
RP/0/RP0/CPU0:router (config-bgp) # session-group group1
RP/0/RP0/CPU0:router (config-bgp-sngrp) # keychain keychain_A
RP/0/RP0/CPU0:router (config-bgp-sngrp) # exit
RP/0/RP0/CPU0:router (config-bgp) # neighbor 172.20.1.1
RP/0/RP0/CPU0:router (config-bgp-nbr) # remote-as 2
RP/0/RP0/CPU0:router (config-bgp-nbr) # use session-group group1
RP/0/RP0/CPU0:router (config-bgp-nbr) # keychain inheritance-disable
```

Related Commands

| Command | Description |
|---------------------------------------|---------------------------------------------------------------------------------|
| keychain, on page 124 | Enables key chain authentication on a TCP connection between two BGP neighbors. |

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

| | |
|----------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>as-number</i> | 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) in 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
Neighbor group configuration
Session group configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| 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.

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.

| Task ID | Task ID | Operations |
|---------|---------|----------------|
| | bgp | read, write |

Examples

The following example shows BGP using autonomous system 30 for the purpose of peering with neighbor 172.20.1.1:

```
RP/0/RP0/CPU0:router(config)# router bgp 140
RP/0/RP0/CPU0:router(config-bgp)# neighbor 172.20.1.1
RP/0/RP0/CPU0:router(config-bgp-nbr)# remote-as 300
RP/0/RP0/CPU0:router(config-bgp-nbr)# local-as 30
```

Related Commands

| Command | Description |
|---------------------------------------------|------------------------------------------------------------------------|
| neighbor-group, on page 141 | Creates a neighbor group and enters neighbor group configuration mode. |
| session-group, on page 189 | Creates a session group and enters session group configuration mode. |

maximum-paths (BGP)

To control the maximum number of parallel routes that Border Gateway Protocol (BGP) installs in the routing table, use the **maximum-paths** command in an appropriate configuration mode. To set the maximum number of parallel routes the software installs to the default value, use the **no** form of this command.

```
maximum-paths {ebgp | ibgp | eibgp} maximum [{unequal-cost}] [{selective}]
no maximum-paths {ebgp | ibgp | eibgp} [{maximum}] [{unequal-cost}]
```

| Syntax Description | |
|---------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| ebgp | Specifies external BGP multipath peers. |
| ibgp | Specifies internal BGP multipath peers. |
| eibgp | Specifies internal and external BGP multipath peers. eiBGP allows simultaneous use of internal and external paths. |
| <i>maximum</i> | Maximum number of parallel routes that BGP installs in the routing table. Range is 2 to 32. |
| unequal-cost | (Optional) Allows iBGP multipaths to have different BGP next-hop Interior Gateway Protocol (IGP) metrics. This option is available when the ibgp keyword is used. |
| selective | (Optional) Allows BGP to be configured such that only routes from selected neighbors can be considered for multipath. This option is used with the multipath option. |

Command Default One path is installed in the routing table.

Command Modes IPv4 address family configuration
IPv6 address family configuration

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | 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.

Use the **maximum-paths** command to allow the BGP to allow the BGP protocol to install multiple paths into the routing table for each prefix. With the eBGP option, 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 with the iBGP option, multiple paths are installed for internal peers that are equal cost based on the BGP best-path algorithm. With the eiBGP option, multiple paths from both iBGP and eBGP are eligible for multipath selection. 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. The **selective** option restricts multipath eligible routes to those that come from peers configured with the **multipath** option.

See *Implementing BGP* in the *Routing Configuration Guide for Cisco NCS 6000 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 super set of the **ibgp** or **ebgp** keywords.

Task ID**Task ID** **Operations**

| Task ID | Operations |
|---------|----------------|
| bgp | read, write |

Examples

The following example shows how to allow a maximum of four paths to a destination to be installed into the IPv4 unicast routing table:

```
RP/0/RP0/CPU0:router(config)# router bgp 109
RP/0/RP0/CPU0:router(config-bgp)# address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-bgp-af)# maximum-paths ebgp 4
RP/0/RP0/CPU0:routerconfig-bgp-af)# commit
```

The following example shows how you can configure selective multipath for iBGP and eBGP peers.



Note This configuration requires the **multipath** option to be configured for the neighbors. See the **multipath** command in the *Cisco ASR 9000 Series Aggregation Services Router Routing Command Reference Guide* for more information.

For information on how this configuration is used, see the BGP Selective Multipath section in the *Cisco ASR 9000 Series Aggregation Services Router Routing Configuration Guide*.

```
RP/0/RP0/CPU0:router(config)# router bgp 1
RP/0/RP0/CPU0:router(config-bgp)# address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-bgp-af)# maximum-paths ibgp 4 selective
RP/0/RP0/CPU0:router(config-bgp-af)# maximum-paths ebgp 5 selective
RP/0/RP0/CPU0:router(config-bgp-af)# commit
```

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

| | |
|-------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>maximum</i> | Maximum number of prefixes allowed from this neighbor. Range is from 1 to 4294967295. |
| | Note When using additional-paths feature, each path with a unique path ID received from a peer is counted separately for the purpose of maximum-prefix functionality. Hence, the <i>maximum</i> value should be configured appropriately when the peer is capable of sending additional-paths. |
| <i>threshold</i> | (Optional) Integer specifying at what percentage of the <i>maximum</i> argument value the software starts to generate a warning message. Range is from 1 to 100. |
| warning-only | (Optional) Instructs the software to only generate a log message when the <i>maximum</i> argument value is exceeded, and not to terminate the peering. |
| restart <i>time-interval</i> | (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.

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 5.2.1 | This command was introduced. |

Usage Guidelines

Use the **maximum-prefix** command to configure a maximum number of prefixes that a BGP router is allowed to receive from a neighbor. It adds another mechanism (besides routing policy) to control prefixes received from a peer.

When the number of received prefixes exceeds the maximum number configured, the software terminates the peering, by default, after sending a cease notification to the neighbor. However, if the **warning-only** keyword is configured, the software writes only a log message, but continues peering with the sender. If the peer is terminated, the peer stays down until the **clear bgp** command is issued or the **restart time-interval** option is used.

This command takes effect immediately if configured on an established neighbor, unless the number of prefixes received from the neighbor already exceeds the configured limits.

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

Task ID

| Task ID | Task | Operations |
|---------|------|----------------|
| | bgp | read, write |

Examples

This example shows the maximum number of IP Version 6 (IPv6) unicast prefixes allowed from neighbor 192.168.40.25 set to 5000, threshold value 80%, and restart time interval 20 minutes:

```
RP/0/RP0/CPU0:router (config) #router bgp 100
RP/0/RP0/CPU0:router (config-bgp) #neighbor 192.168.40.25
RP/0/RP0/CPU0:router (config-bgp-nbr) #remote-as 1
RP/0/RP0/CPU0:router (config-bgp-nbr) #address-family ipv6 unicast
RP/0/RP0/CPU0:router (config-bgp-nbr-af) #maximum-prefix 5000 80 restart 20
```

This example shows the maximum number of IP Version 4 (IPv4) unicast prefixes allowed from the neighbor 192.168.40.24 set to 1000:

```
RP/0/RP0/CPU0:router (config-bgp) # router bgp 109
RP/0/RP0/CPU0:router (config-bgp) # neighbor 192.168.40.24
RP/0/RP0/CPU0:router (config-bgp-nbr) # remote-as 1
RP/0/RP0/CPU0:router (config-bgp-nbr) #address-family ipv4 unicast
RP/0/RP0/CPU0:router (config-bgp-nbr-af) # maximum-prefix 1000
```

Related Commands

| Command | Description |
|---------------------------------------------|-------------------------------------------------------------------------------------------------------|
| af-group, on page 10 | Creates an address family group for BGP neighbors and enters address family group configuration mode. |
| clear bgp, on page 82 | Resets a BGP connection using BGP hard or soft reconfiguration. |
| neighbor-group, on page 141 | 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

Command History

| Release | Modification |
|---------------|------------------------------|
| 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.

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 |

Related Commands

| Command | Description |
|-------------------------------------------------|--------------------------------------------------------------------------------|
| address-family (BGP), on page 6 | 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/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router bgp 1
RP/0/RP0/CPU0:router(config-bgp)# address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-bgp-af)# maximum-paths ibgp 4 selective
RP/0/RP0/CPU0:router(config-bgp-af)# maximum-paths ebgp 5 selective
RP/0/RP0/CPU0:router(config-bgp-af)# neighbor 1.1.1.1
RP/0/RP0/CPU0:router(config-bgp-nbr)# address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-bgp-nbr-af)# multipath
RP/0/RP0/CPU0:router(config-bgp-nbr-af)# commit
```

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 | <i>ip-address</i> IPv4 or IPv6 IP address of the BGP-speaking neighbor. | | | | |
|---------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------|--------------|---------------|------------------------------|
| Command Default | Neighbor mode is not specified. | | | | |
| Command Modes | Router configuration | | | | |
| Command History | <table border="1"> <thead> <tr> <th>Release</th> <th>Modification</th> </tr> </thead> <tbody> <tr> <td>Release 5.0.0</td> <td>This command was introduced.</td> </tr> </tbody> </table> | Release | Modification | Release 5.0.0 | This command was introduced. |
| Release | Modification | | | | |
| 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.

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 a remote autonomous system number, and an IP address and address family must be enabled on the neighbor.

Unlike IPv4, IPv6 must be enabled before any IPv6 neighbors can be defined. Enable IPv6 in router configuration mode using the **address-family** command.



Note Configuration for the neighbor cannot occur (peering is not established) until the neighbor is given a remote as-number and neighbor address.

The **no** form of this command causes the peering with the neighbor to be terminated and all configuration that relates to the neighbor to be removed.

| Task ID | Task ID | Operations |
|---------|---------|----------------|
| | bgp | read, write |

Examples

The following example shows how to place the router in neighbor configuration mode for BGP routing process 1 and configure the neighbor IP address 172.168.40.24 as a BGP peer:

```
RP/0/RP0/CPU0:router(config)# router bgp 1
RP/0/RP0/CPU0:router(config-bgp)# neighbor 172.168.40.24
RP/0/RP0/CPU0:router(config-bgp-nbr)# remote-as 65000
```

The following example shows how to enable IPv6 for BGP, then place the router in neighbor configuration mode for an IPv6 neighbor, 3000::1, and configure neighbor 3000::1 as a BGP peer:

```
RP/0/RP0/CPU0:router(config)# router bgp 100
RP/0/RP0/CPU0:router(config-bgp)# address-family ipv6 unicast
RP/0/RP0/CPU0:router(config-bgp-af)# exit
RP/0/RP0/CPU0:router(config-bgp)# neighbor 3000::1
RP/0/RP0/CPU0:router(config-bgp-nbr)# remote-as 2002
RP/0/RP0/CPU0:router(config-bgp-nbr)# address-family ipv6 unicast
```

Related Commands

| Command | Description |
|-------------------------------------------------|-----------------------------------------------------------------------------------------|
| address-family (BGP), on page 6 | Enters address family configuration mode for configuring BGP routing sessions. |
| remote-as (BGP), on page 173 | Adds an entry to the BGP neighbor table. |
| use, on page 368 | 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 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.

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

Related Commands

| Command | Description |
|-------------------------------------------------|----------------------------------------------------------------------------------------------|
| address-family (BGP), on page 6 | Enters various address family configuration modes for configuring BGP routing sessions. |
| neighbor (BGP), on page 139 | Enters neighbor configuration mode for configuring BGP routing sessions. |
| use, on page 368 | Inherits characteristics from a neighbor group, a session group, or an address family group. |

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 | | |
|----------------------------------------------|------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>ip-address</i> | | Network that BGP advertises. |
| <i>/prefix-length</i> | | 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. |
| <i>ip-address mask</i> | | Network mask applied to the <i>ip-address</i> argument. |
| route-policy <i>route-policy-name</i> | (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

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | 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.

A network specified with this command is originated and advertised to neighbors only if there exists a route for the network in the routing table. That is, there must be a route learned using local or connected networks, static routing, or a dynamic IGP such as IS-IS or OSPF.

Other than the available system resources on the router, no limit exists on the number of network commands that can be configured.

| Task ID | Task ID | Operations |
|---------|---------|----------------|
| | bgp | read, write |

Examples

The following example shows how to configure the local router to originate the IPv4 unicast network 172.20.0.0/16:

```
RP/0/RP0/CPU0:router(config)#router bgp 120  
RP/0/RP0/CPU0:router(config-bgp)# address-family ipv4 unicast  
RP/0/RP0/CPU0:router(config-bgp-af)# network 172.20.0.0/16
```

Related Commands

| Command | Description |
|-------------------------------------------------|-------------------------------------------------------------------------------------------------------|
| network backdoor, on page 145 | Specifies a backdoor route to a BGP border router that provides better information about the network. |
| redistribute (BGP), on page 169 | Redistributes routes from one routing domain into another routing domain. |

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

| | |
|-----------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>ip-address</i> | Network that provides a backdoor route. |
| <i>/prefix-length</i> | 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. |
| <i>mask</i> | Network mask applied to the <i>ip-address</i> argument. |

Command Default

No backdoor routes are installed.

Command Modes

IPv4 address family configuration
 IPv6 address family configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| 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.

Configuring the **network backdoor** command does not cause BGP to originate a network, even if an IGP route for the network exists. Ordinarily, the backdoor network would be learned through both an eBGP and IGP. The BGP best-path selection algorithm does not change when a network is configured as a backdoor network.

Task ID

| Task ID | Operations |
|---------|----------------|
| bgp | read, write |

Examples

The following example shows IP Version 4 (IPv4) unicast network 192.168.40.0/24 configured as a backdoor network:

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

```
RP/0/RP0/CPU0:router(config-bgp-af)# network 192.168.40.0/24 backdoor
```

Related Commands

| Command | Description |
|--------------------------------------------|---------------------------------------------------------------------------------------------------------|
| network (BGP), on page 143 | Specifies a local network that the BGP routing process should originate and advertise to its neighbors. |

next-hop-self

To disable next-hop calculation and insert your own address in the next-hop field of Border Gateway Protocol (BGP) updates, use the **next-hop-self** command in an appropriate configuration mode. To enable next-hop calculation, use the **no** form of this command.

```
next-hop-self [{inheritance-disable}]
no next-hop-self [{inheritance-disable}]
```

| | |
|---------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax Description | inheritance-disable (Optional) Allows a next-hop calculation override when this feature may be inherited from a neighbor group or address family group. |
|---------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------|

| | |
|------------------------|------------------------------------------------------------------------------------------------------------------|
| Command Default | When this command is not specified, the software calculates the next hop for BGP updates accepted by the router. |
|------------------------|------------------------------------------------------------------------------------------------------------------|

| | |
|----------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Command Modes | IPv4 address family group configuration IPv6 address family group configuration IPv4 neighbor address family configuration IPv4 neighbor group address family configuration IPv6 neighbor group address family configuration IPv4 labeled-unicast address family configuration IPv6 labeled-unicast address family configuration |
|----------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

| Command History | <table border="1"> <thead> <tr> <th>Release</th> <th>Modification</th> </tr> </thead> <tbody> <tr> <td>Release 5.0.0</td> <td>This command was introduced.</td> </tr> </tbody> </table> | Release | Modification | Release 5.0.0 | This command was introduced. |
|------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------|--------------|---------------|------------------------------|
| Release | Modification | | | | |
| 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. |
|-------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

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.

| Task ID | <table border="1"> <thead> <tr> <th>Task ID</th> <th>Operations</th> </tr> </thead> <tbody> <tr> <td>bgp</td> <td>read, write</td> </tr> </tbody> </table> | Task ID | Operations | bgp | read, write |
|----------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------|---------|------------|-----|----------------|
| Task ID | Operations | | | | |
| bgp | read, write | | | | |

Examples

The following example shows how to set the next hop of the update field for all IP Version 4 (IPv4) unicast routes advertised to neighbor 172.20.1.1 to an address of the local router:

```
RP/0/RP0/CPU0:router(config)# router bgp 140
RP/0/RP0/CPU0:router(config-bgp)# neighbor 172.20.1.1
RP/0/RP0/CPU0:router(config-bgp-nbr)# remote-as 1
RP/0/RP0/CPU0:router(config-bgp-nbr)# address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-bgp-nbr-af)# next-hop-self
```

The following example shows how to disable the **next-hop-self** command for neighbor 172.20.1.1. If not overridden, the next hop would be inherited from address family group group1:

```
RP/0/RP0/CPU0:router(config)# router bgp 140
RP/0/RP0/CPU0:router(config-bgp)# af-group group1 address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-bgp-afgrp)# next-hop-self
RP/0/RP0/CPU0:router(config-bgp-afgrp)# exit
RP/0/RP0/CPU0:router(config-bgp)# neighbor 172.20.1.1
RP/0/RP0/CPU0:router(config-bgp-nbr)# remote-as 1
RP/0/RP0/CPU0:router(config-bgp-nbr)# address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-bgp-nbr-af)# use af-group group1
RP/0/RP0/CPU0:router(config-bgp-nbr-af)# next-hop-self inheritance-disable
```

Related Commands

| Command | Description |
|---------------------------------------------|-------------------------------------------------------------------------------------------------------|
| af-group, on page 10 | Creates an address family group for BGP neighbors and enters address family group configuration mode. |
| neighbor-group, on page 141 | Creates a neighbor group and enters neighbor group configuration mode. |
| use, on page 368 | 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 | 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 5.0.0 | This command was introduced. |

| | |
|-------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Usage Guidelines | Use the next-hop-unchanged command to propagate the next hop unchanged for multihop eBGP peering sessions. This command should not be configured on a route reflector, and the next-hop-self command should not be used to modify the next-hop attribute for a route reflector when this feature is enabled for a route reflector client. |
|-------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|



| | |
|-------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Note | Incorrectly setting BGP attributes for a route reflector can cause inconsistent routing, routing loops, or a loss of connectivity. Setting BGP attributes for a route reflector should be attempted only by an experienced network operator. |
|-------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

| | | |
|----------------|----------------|-------------------|
| Task ID | Task ID | Operations |
| | bgp | read, write |

Examples

The following example shows how to disable the overwriting of next hops before advertising to eBGP peers:

```
RP/0/RP0/CPU0:router(config)# router bgp 140
```

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

Related Commands

| Command | Description |
|--------------------------------------------|---------------------------------------------------------------------------------------------------------------|
| next-hop-self, on page 147 | Disables next-hop calculation and allows you to insert your own address in the next-hop field of BGP updates. |
| use, on page 368 | Inherits characteristics from a neighbor group, session group, or address family group. |

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 | <i>route-policy-name</i> 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 | | | | |
| Command History | <table border="1"> <thead> <tr> <th>Release</th> <th>Modification</th> </tr> </thead> <tbody> <tr> <td>Release 5.0.0</td> <td>This command was introduced.</td> </tr> </tbody> </table> | Release | Modification | Release 5.0.0 | This command was introduced. |
| Release | Modification | | | | |
| Release 5.0.0 | This command was introduced. | | | | |
| Usage Guidelines | <p>To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.</p> <p>Use the nexthop route-policy command to configure route policy filtering using next hops.</p> <p>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:</p> <ul style="list-style-type: none"> • 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. <p>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.</p> <p>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.</p> | | | | |
| Task ID | <table border="1"> <thead> <tr> <th>Task ID</th> <th>Operations</th> </tr> </thead> <tbody> <tr> <td>bgp</td> <td>read, write</td> </tr> </tbody> </table> | Task ID | Operations | bgp | read, write |
| 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:

```
RP/0/RP0/CPU0:router(config)# router bgp 109  
RP/0/RP0/CPU0:router(config-bgp)# address-family ipv4 unicast  
RP/0/RP0/CPU0:router(config-bgp-af)# nexthop route-policy nexthop_A
```

Related Commands

| Command | Description |
|------------------------------------------------|--------------------------------------------------------------------|
| <code>route-policy (RPL)</code> | Defines a route policy and enters route-policy configuration mode. |
| show bgp nexthops, on page 275 | 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 | non-critical |
|--------------------|------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------|
| | Specifies critical next-hop events. For example, when the next hop is unreachable. | |
| | <i>delay</i> | Trigger delay, in milliseconds. Range is 0 to 4294967295. |
| | | Specifies noncritical next-hop events. For example, Interior Gateway Protocol (IGP) metric changes. |

| Command Default |
|--------------------------------------------------------------------------------|
| critical : 3000 msec for IPv4 address family and IPv6 address family |
| critical : 0 msec for VPNv4 address family and VPNv6 address family |
| non-critical : 10000 msec IPv4, IPv6, VPNv4, and VPNv6 address families |

| Command Modes |
|-----------------------------------|
| IPv4 address family configuration |
| Pv6 address family configuration |

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | 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.

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).

| Task ID | Task ID | Operations |
|---------|---------|----------------|
| | bgp | read, write |

Examples

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

```
RP/0/RP0/CPU0:router(config)# router bgp 109  
RP/0/RP0/CPU0:router(config-bgp)# address-family ipv4 unicast  
RP/0/RP0/CPU0:router(config-bgp-af)# nexthop trigger-delay critical 3500
```

nsr (BGP)

To activate Border Gateway Protocol (BGP) nonstop routing (NSR), use the **nsr** command in BGP XR Config mode. To deactivate BGP NSR, use the **no nsr** 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 XR Config

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | Release 5.2.1 | 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 **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/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router (config)# router bgp 120
RP/0/RP0/CPU0:router (config-bgp)# nsr
```

The following example shows how to disable BGP NSR:

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router (config)# router bgp 120
RP/0/RP0/CPU0:router (config-bgp)# no nsr
```

Related Commands

| Command | Description |
|-------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| router bgp, on page 182 | 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 277 | 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 XR Config mode. To re-enable BGP NSR, use the **no nsr disable** 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 XR Config

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | Release 5.2.3 | 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/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router bgp 120
RP/0/RP0/CPU0:router(config-bgp)# nsr disable
```

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

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router bgp 120
RP/0/RP0/CPU0:router(config-bgp)# no nsr disable
```

Related Commands

| Command | Description |
|-------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| router bgp, on page 182 | 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 277 | Displays Border Gateway Protocol (BGP) nonstop routing (NSR) information. |

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 | <i>route-policy-name</i> 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 |
|----------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

| Command History | Release | Modification |
|------------------------|----------------|------------------------------|
| | 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 | Operations |
|----------------|----------------|-------------------|
| | bgp | read, write |

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

```
RP/0/RP0/CPU0:router(config)#router bgp 6
RP/0/RP0/CPU0:router(config-bgp)# neighbor 172.20.1.1
RP/0/RP0/CPU0:router(config-bgp-nbr)# remote-as 1
RP/0/RP0/CPU0:router(config-bgp-nbr)# address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-bgp-nbr-af)#orf route-policy policy_A
```

Related Commands

| Command | Description |
|-------------------------------------------------|------------------------------------------------------------------------------------|
| route-policy (BGP), on page 178 | 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 | <p>clear Specifies that an unencrypted password follows. The password must be a case-sensitive, clear-text unencrypted password.</p> <p>encrypted Specifies that an encrypted password follows. The password must be a case-sensitive, encrypted password.</p> <p><i>password</i> 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."</p> | | | | |
|---------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------|--------------|---------------|------------------------------|
| 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 | <p>Neighbor configuration</p> <p>Neighbor group configuration</p> <p>Session group configuration</p> | | | | |
| Command History | <table border="1"> <thead> <tr> <th data-bbox="386 1171 938 1203">Release</th> <th data-bbox="959 1171 1526 1203">Modification</th> </tr> </thead> <tbody> <tr> <td data-bbox="386 1224 938 1255">Release 5.0.0</td> <td data-bbox="959 1224 1526 1255">This command was introduced.</td> </tr> </tbody> </table> | Release | Modification | Release 5.0.0 | This command was introduced. |
| Release | Modification | | | | |
| Release 5.0.0 | This command was introduced. | | | | |
| Usage Guidelines | <p>To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.</p> <p>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.</p> <p>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.</p> <p>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.</p> | | | | |

| Task ID | Task ID | Operations |
|---------|---------|----------------|
| | bgp | read, write |

Examples

The following example shows how to configure neighbor 172.20.1.1 to use MD5 authentication with the password password1:

```
RP/0/RP0/CPU0:router(config)# router bgp 140
RP/0/RP0/CPU0:router(config-bgp)#neighbor 172.20.1.1
RP/0/RP0/CPU0:router(config-bgp-nbr)#remote-as 1
RP/0/RP0/CPU0:router(config-bgp-nbr)#password clear password1
```

Related Commands

| Command | Description |
|-----------------------------------------------|----------------------------------------------------------------------------------------------------------|
| neighbor-group, on page 141 | Creates a neighbor group and enters neighbor group configuration mode. |
| password-disable, on page 163 | Overrides any inherited password configuration from a neighbor group or session group for BGP neighbors. |
| session-group, on page 189 | Creates a session group and enters session group configuration mode. |
| timers (BGP), on page 351 | 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 Neighbor group configuration Session group configuration | |
| Command History | Release | Modification |
| | 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 you specify a password on a neighbor group or session group, all users of the group inherit the password. Specifying a different **password** command specifically on a neighbor that uses the group overrides the inherited value. Specifying **password-disable** on a neighbor that uses the group disables password authentication for the neighbor.

| Task ID | Task ID | Operations |
|---------|---------|----------------|
| | bgp | read, write |

Examples

The following example shows how to disable MD5 authentication for neighbor 172.20.1.1, preventing it from inheriting the password password1 from session group group1:

```
RP/0/RP0/CPU0:router(config)# router bgp 140
RP/0/RP0/CPU0:router(config-bgp)# session-group group1
RP/0/RP0/CPU0:router(config-bgp-sngrp)# password clear password1
RP/0/RP0/CPU0:router(config-bgp-sngrp)# exit
RP/0/RP0/CPU0:router(config-bgp)# neighbor 172.20.1.1
RP/0/RP0/CPU0:router(config-bgp-nbr)# remote-as 2
RP/0/RP0/CPU0:router(config-bgp-nbr)# use session-group group1
RP/0/RP0/CPU0:router(config-bgp-nbr)# password-disable
```

Related Commands

| Command | Description |
|---------------------------------------------|----------------------------------------------------------------------------------------------|
| neighbor-group, on page 141 | Creates a neighbor group and enters neighbor group configuration mode. |
| password (BGP), on page 161 | Enables MD5 authentication on a TCP connection between two BGP neighbors. |
| session-group, on page 189 | Creates a session group and enters session group configuration mode. |
| use, on page 368 | Inherits characteristics from a neighbor group, a session group, or an address family group. |

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 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.

Use the **precedence** command to set the precedence value.

Task ID

| Task ID | Operations |
|---------|----------------|
| bgp | read, write |

Examples

The following example shows how to set the precedence to 2:

```
RP/0/RP0/CPU0:router(config)# router bgp 5  
RP/0/RP0/CPU0:router(config-bgp)# neighbor 10.1.1.1  
RP/0/RP0/CPU0:router(config-bgp-nbr)# remote-as 100  
RP/0/RP0/CPU0:router(config-bgp-nbr)# precedence 2
```

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

Neighbor group configuration

Session group configuration

Command History

Release

Release 5.0.0

Modification

This command was introduced.

Usage Guidelines

To use this command, you must be in a user group associated with a task group that 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/RP0/CPU0:router(config)# router bgp 1
RP/0/RP0/CPU0:router(config-bgp)# neighbor 172.20.1.1
RP/0/RP0/CPU0:router(config-bgp-nbr)# remote-as 1
RP/0/RP0/CPU0:router(config-bgp-nbr)# receive-buffer-size 65536 8192
```

Related Commands

| Command | Description |
|---------------------------------------------------------|------------------------------------------------------------------------|
| neighbor-group, on page 141 | Creates a neighbor group and enters neighbor group configuration mode. |
| send-buffer-size, on page 183 | Sets the size of the send buffers for a BGP neighbor. |
| session-group, on page 189 | Creates a session group and enters session group configuration mode. |
| socket receive-buffer-size, on page 341 | 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]
```

| Syntax Description | | |
|----------------------------------------------|--|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| connected | | Redistributes connected routes. Connected routes are established automatically when IP is enabled on an interface. |
| metric <i>metric-value</i> | | (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 <i>route-policy-name</i> | | (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 address family configuration mode. |

| | |
|---------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>process-id</i> | <p>For the eigrp keyword, an EIGRP instance name from which routes are to be redistributed.</p> <p>For the isis keyword, an IS-IS instance name from which routes are to be redistributed.</p> <p>For the ospf keyword, an OSPF instance name from which routes are to be redistributed.</p> <p>The <i>process-id</i> value takes the form of a string. A decimal number can be entered, but it is stored internally as a string.</p> |
| match { internal external [1 2] nssa-external [1 2] } | <p>(Optional) Specifies the criteria by which OSPF routes are redistributed into other routing domains. It can be one or more of the following:</p> <ul style="list-style-type: none"> • 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. <p>For the external and nssa-external options, if a type is not specified, then both Type 1 and Type 2 are assumed.</p> |
| isis | <p>Specifies that routes are distributed from the IS-IS protocol.</p> <p>Redistribution from IS-IS is allowed under IPv4 unicast, IPv6 unicast, and address-families.</p> |
| level { 1 1-inter-area 2 } | <p>(Optional) Specifies the IS-IS level from which routes are redistributed. It can be one of the following:</p> <ul style="list-style-type: none"> • 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 | <p>Specifies that routes are distributed from the OSPF protocol. You must be in IPv4 unicast address family configuration mode.</p> |
| rip | <p>Specifies that routes are distributed from RIP. You must be in IPv4 unicast address family configuration mode.</p> |
| static | <p>Redistributes IP static routes.</p> |

Command Default

Route redistribution is disabled.

For IS-IS, the default is to redistribute Level 1 and Level 2 routes.

For OSPF, the default is to redistribute internal, external, and NSSA external routes of Type 1 and Type 2.

For OSPFv3, the default is to redistribute internal, external, and NSSA external routes of Type 1 and Type 2.

By default, the Interior Gateway Protocol (IGP) metric is assigned to the route. For connected and static routes the default metric is 0.

metric *metric-value*: 0

match { **internal** | **external** [1 | 2] | **nssa-external** [1 | 2] }; If no match is specified, the default is to match all routes.

Command Modes IPv4 address family configuration (**connected**, **eigrp**, **isis**, **ospf**, **rip**, and **static** are supported)

IPv6 address family configuration (**connected**, **eigrp**, **isis**,

ospfv3,
and **static** are supported)

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | 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.



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.

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.

| Task ID | Task ID | Operations |
|---------|---------|----------------|
| | bgp | read, write |

Examples The following example shows how to redistribute IP Version 4 (IPv4) unicast OSPF routes from OSPF instance 110 into BGP:

```
RP/0/RP0/CPU0:router(config)# router bgp 109
RP/0/RP0/CPU0:router(config-bgp)# address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-bgp-af)# redistribute ospf 110
```

Related Commands

| Command | Description |
|--------------------------------------------|---------------------------------------------------------------------------------------------------------|
| network (BGP), on page 143 | Specifies a local network that the BGP routing process should originate and advertise to its neighbors. |

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) in asdot format is 1.0 to 65535.65535.

Command Default

No BGP neighbors exist.

Command Modes

Neighbor configuration
 Neighbor group configuration
 Session group configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| 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.

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 Operations ID

bgp read,
write

Examples

The following example shows how to assign autonomous system numbers on two neighbors, neighbor 10.0.0.1, (internal) and neighbor 192.168.0.1 (external), setting up a peering session that shares routing information between this router and each of these neighbors:

```
RP/0/RP0/CPU0:router(config)# router bgp 1
RP/0/RP0/CPU0:router(config-bgp)# session-group group2
RP/0/RP0/CPU0:router(config-bgp-sngrp)# remote-as 1
RP/0/RP0/CPU0:router(config-bgp-sngrp)# exit
RP/0/RP0/CPU0:router(config-bgp)# neighbor 10.0.0.1
RP/0/RP0/CPU0:router(config-bgp-nbr)# use session-group group2
```

The following example shows how to configure a session group called group2 with an autonomous system number 1. Neighbor 10.0.0.1 is created when it inherits the autonomous system number 1 from session group group2.

```
RP/0/RP0/CPU0:router(config)#router bgp 1
RP/0/RP0/CPU0:router(config-bgp)# session-group group2
RP/0/RP0/CPU0:router(config-bgp-sngrp)# remote-as 1
RP/0/RP0/CPU0:router(config-bgp-sngrp)# exit
RP/0/RP0/CPU0:router(config-bgp)# neighbor 10.0.0.1
RP/0/RP0/CPU0:router(config-bgp-nbr)# use session-group group2
```

Related Commands

| Command | Description |
|---------------------------------------------|--------------------------------------------------------------------------|
| neighbor (BGP), on page 139 | Enters neighbor configuration mode for configuring BGP routing sessions. |
| neighbor-group, on page 141 | Creates a neighbor group and enters neighbor group configuration mode. |
| router bgp, on page 182 | Configures the BGP routing process. |
| session-group, on page 189 | Creates a session group and enters session group configuration mode. |

| Command | Description |
|----------------------------------|-----------------------------------------------------------------------------------------|
| use, on page 368 | 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 |

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | 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. |

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/RP0/CPU0:router(config)# router bgp 140
RP/0/RP0/CPU0:router(config-bgp)# neighbor 172.20.1.1
RP/0/RP0/CPU0:router(config-bgp-nbr)# remote-as 1
RP/0/RP0/CPU0:router(config-bgp-nbr)# address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-bgp-nbr-af)# remove-private-as
```

The following example shows how to disable the remove private autonomous system number feature for neighbor 172.20.1.1, preventing this feature from being automatically inherited from address family group group1:

```
RP/0/RP0/CPU0:router(config)# router bgp 140
RP/0/RP0/CPU0:router(config-bgp)# af-group group1 address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-bgp-afgrp)# remove-private-as
RP/0/RP0/CPU0:router(config-bgp-afgrp)# exit
RP/0/RP0/CPU0:router(config-bgp)# neighbor 172.20.1.1
RP/0/RP0/CPU0:router(config-bgp-nbr)# remote-as 1
RP/0/RP0/CPU0:router(config-bgp-nbr)# address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-bgp-nbr-af)# use af-group group1
RP/0/RP0/CPU0:router(config-bgp-nbr-af)# remove-private-as inheritance-disable
```

Related Commands

| Command | Description |
|----------------------------------------------|-------------------------------------------------------------------------------------------------------|
| af-group, on page 10 | Creates an address family group for BGP neighbors and enters address family group configuration mode. |
| neighbor-group, on page 141 | Creates a neighbor group and enters neighbor group configuration mode. |
| remote-as (BGP), on page 173 | Allows entries to the BGP neighbor table. |

route-policy (BGP)

To apply a routing policy to updates advertised to or received from a Border Gateway Protocol (BGP) neighbor, use the **route-policy** command in an appropriate configuration mode. To disable applying routing policy to updates, use the **no** form of this command.

```
route-policy route-policy-name [{parameter1, parameter2, . . . , parametern}] {in | out}
no route-policy route-policy-name [{parameter1, parameter2, . . . , parametern}] {in | out}
```

| Syntax Description | |
|--------------------------|---------------------------------------------------------------------------------------------------------|
| <i>route-policy-name</i> | Name of route policy. Up to 16 parameters can follow the route-policy-name, enclosed in brackets ([]). |
| in | Applies policy to inbound routes. |
| out | Applies policy to outbound routes. |

Command Default No policy is applied.

Command Modes

- IPv4 address family group configuration
- IPv6 address family group configuration
- IPv4 neighbor address family configuration
- IPv4 neighbor group address family configuration
- IPv6 neighbor group address family configuration

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | 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.

Use the **route-policy** command to specify a routing policy for an inbound or outbound route. The policy can be used to filter routes or modify route attributes. The **route-policy** command is used to define a policy.



Note Configuring a large number of uniquely named outbound neighbor policies can adversely affect performance. This is true even if the uniquely named route policies are functionally identical. The user is discouraged from configuring multiple functionally identical route policies for use with this command. For example, if Policy A and Policy B are identical but named for different neighbors, the two policies should be configured as a single policy.

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

| Task ID | Task ID | Operations |
|---------|---------|----------------|
| | bgp | read, write |

Examples

The following example shows how to apply the In-Ipv4 policy to inbound IP Version 4 (IPv4) unicast routes from neighbor 172.20.1.1:

```
RP/0/RP0/CPU0:router(config)# router bgp 1
RP/0/RP0/CPU0:router(config-bgp)# neighbor 172.20.1.1
RP/0/RP0/CPU0:router(config-bgp-nbr)# remote-as 1
RP/0/RP0/CPU0:router(config-bgp-nbr)# address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-bgp-nbr-af)# route-policy In-Ipv4 in
```

Related Commands

| Command | Description |
|---------------------------------------------|-------------------------------------------------------------------------------------------------------|
| af-group, on page 10 | Creates an address family group for BGP neighbors and enters address family group configuration mode. |
| neighbor-group, on page 141 | 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 | |
| Command History | Release | Modification |
| | Release 5.0.0 | This command was introduced. |
| Usage Guidelines | <p>To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.</p> <p>This command is restricted to internal BGP (iBGP) neighbors only.</p> <p>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.</p> <p>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.</p> <p>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.</p> <p>The local router is a route reflector as long as it has at least one route reflector client.</p> <p>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.</p> | |

| 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/RP0/CPU0:router(config)# router bgp 140
RP/0/RP0/CPU0:router(config-bgp)# neighbor 172.20.1.1
RP/0/RP0/CPU0:router(config-bgp-nbr)# remote-as 140
RP/0/RP0/CPU0:router(config-bgp-nbr)# address-family ipv4 unicast
RP/0/RP0/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 group1:

```
RP/0/RP0/CPU0:router(config)# router bgp 140
RP/0/RP0/CPU0:router(config-bgp)# af-group group1 address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-bgp-afgrp)# route-reflector-client
RP/0/RP0/CPU0:router(config-bgp-afgrp)#exit
RP/0/RP0/CPU0:router(config-bgp)# neighbor 172.20.1.1
RP/0/RP0/CPU0:router(config-bgp-nbr)# remote-as 140
RP/0/RP0/CPU0:router(config-bgp-nbr)# address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-bgp-nbr-af)# use af-group group1
RP/0/RP0/CPU0:router(config-bgp-nbr-af)# route-reflector-client inheritance-disable
```

Related Commands

| Command | Description |
|---------------------------------------------|-------------------------------------------------------------------------------------------------------|
| af-group, on page 10 | Creates an address family group for BGP neighbors and enters address family group configuration mode. |
| bgp cluster-id, on page 46 | Configures the cluster ID if the BGP cluster has more than one route reflector. |
| neighbor-group, on page 141 | 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 XR Config 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) in asdot format is 1.0 to 65535.65535.

Command Default

No BGP routing process is enabled.

Command Modes

XR Config mode

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 5.0.0 | This command was introduced. |

Usage Guidelines

Use the **router bgp** command to set up a distributed routing core that automatically guarantees the loop-free exchange of routing information between autonomous systems.

Task ID

| Task ID | Operations |
|---------|----------------|
| bgp | read, write |
| rib | read, write |

Examples

The following example shows how to configure a BGP process for autonomous system 120:

```
RP/0/RP0/CPU0:router(config)# router bgp 120
```

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

Neighbor group configuration

Session group configuration

Command History

Release

Release 5.0.0

Modification

This command was introduced.

Usage Guidelines

To use this command, you must be in a user group associated with a task group that 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/RP0/CPU0:router(config)# router bgp 1
RP/0/RP0/CPU0:router(config-bgp)# neighbor 172.20.1.1
RP/0/RP0/CPU0:router(config-bgp-nbr)# remote-as 1
RP/0/RP0/CPU0:router(config-bgp-nbr)# send-buffer-size 8192 8192
```

Related Commands

| Command | Description |
|------------------------------------------------------|------------------------------------------------------------------------|
| neighbor-group, on page 141 | Creates a neighbor group and enters neighbor group configuration mode. |
| receive-buffer-size, on page 167 | Sets the size of the receive buffers for a BGP neighbor. |
| session-group, on page 189 | Creates a session group and enters session group configuration mode. |
| socket send-buffer-size, on page 343 | 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 |
|----------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

| Command History | Release | Modification |
|------------------------|----------------|------------------------------------------|
| | Release 5.0.0 | This command was introduced. |
| | Release 6.3.1 | This command supports large communities. |

| | |
|-------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Usage Guidelines | To use this command, you must be in a user group associated with a task group that 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.

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 |

Related Commands

| Command | Description |
|-----------------------------------------------------------|-------------------------------------------------------------------------------------------------------|
| af-group, on page 10 | Creates an address family group for BGP neighbors and enters address family group configuration mode. |
| neighbor-group, on page 141 | Creates a neighbor group and enters neighbor group configuration mode. |
| send-extended-community-ebgp, on page 187 | Specifies that extended community attributes are sent to eBGP neighbors. |

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 |
|----------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

| | | |
|------------------------|----------------|------------------------------|
| Command History | Release | Modification |
| | 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. |
|-------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

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.

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 |

Related Commands

| Command | Description |
|--------------------------------------------------|-------------------------------------------------------------------------------------------------------|
| af-group, on page 10 | Creates an address family group for BGP neighbors and enters address family group configuration mode. |
| neighbor-group, on page 141 | Creates a neighbor group and enters neighbor group configuration mode. |
| send-community-ebgp, on page 185 | 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 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. |
|-------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

Use the **session-group** command to create a session group from which neighbors can inherit configuration that is address family-independent. That is, session groups cannot have address family-specific configuration. This command enters the session group configuration mode in which configuration for a session group is entered.

Many commands can be configured in both session group configuration mode and neighbor configuration mode.

Use of session groups saves time and reduces the router configuration size. Because the configuration of a session group can be inherited by any number of neighbors, use of the group can eliminate the need to copy long or complex configurations on each of a large number of neighbors. A neighbor can inherit all configuration from a session group simply by configuring the **use** command. Specific inherited session group configuration commands can be overridden for a specific neighbor by explicitly configuring the command for the specific neighbor.

The **no** form of this command causes all of the configuration for the session group to be removed. You cannot use the **no** form of this command if removing the group would leave one or more neighbors without a configured remote autonomous system number.

| Task ID | Task ID | Operations |
|----------------|----------------|-------------------|
| | bgp | read, write |

Examples

The following example shows a session group called `group1` that is used by two neighbors, 10.0.0.1 and 10.0.0.2. Because `group1` is a session group, it contains only address family-independent configuration. And because `group1` is used by neighbors 10.0.0.1 and 10.0.0.2, they inherit the configuration of the group.

```
RP/0/RSP0RP0/CPU0:router(config)# router bgp 1
RP/0/RSP0RP0/CPU0:router(config-bgp)# session-group group1
RP/0/RSP0RP0/CPU0:router(config-bgp-sngrp)# remote-as 1
RP/0/RSP0RP0/CPU0:router(config-bgp-sngrp)# advertisement-interval 2
RP/0/RSP0RP0/CPU0:router(config-bgp-sngrp)# exit
RP/0/RSP0RP0/CPU0:router(config-bgp)# neighbor 10.0.0.1
RP/0/RSP0RP0/CPU0:router(config-bgp-nbr)# use session-group group1
RP/0/RSP0RP0/CPU0:router(config-bgp-nbr)# exit
RP/0/RSP0RP0/CPU0:router(config-bgp)# neighbor 10.0.0.2
RP/0/RSP0RP0/CPU0:router(config-bgp-nbr)# use session-group group1
```

The following example shows a session group called `group1` used by two neighbors, 10.0.0.1 and 10.0.0.2. Because `group1` is a session group, it contains only address family-independent configuration. And because `group1` is used by neighbors 10.0.0.1 and 10.0.0.2, they inherit the configuration of the group. However, the `password password1` configuration from `group1` is overridden for neighbor 10.0.0.2, using the `password-disable` command in the neighbor 10.0.0.2 configuration submode.

```
RP/0/RP0/CPU0:router(config)# router bgp 1
RP/0/RP0/CPU0:router(config-bgp)# session-group group1
RP/0/RP0/CPU0:router(config-bgp-sngrp)# remote-as 1
RP/0/RP0/CPU0:router(config-bgp-sngrp)# advertisement-interval 2
RP/0/RP0/CPU0:router(config-bgp-sngrp)# password password1
RP/0/RP0/CPU0:router(config-bgp-sngrp)# exit
RP/0/RP0/CPU0:router(config-bgp)# neighbor 10.0.0.1
RP/0/RP0/CPU0:router(config-bgp-nbr)# use session-group group1
RP/0/RP0/CPU0:router(config-bgp-nbr)# exit
RP/0/RP0/CPU0:router(config-bgp)# neighbor 10.0.0.2
RP/0/RP0/CPU0:router(config-bgp-nbr)# use session-group group1
RP/0/RP0/CPU0:router(config-bgp-nbr)# password-disable
```

session-open-mode

To establish a Border Gateway Protocol (BGP) session with a specific TCP open mode, use the **session-open-mode** command in an appropriate configuration mode. To restore the default state, use the **no** form of this command.

```
session-open-mode {active-only | both | passive-only}
no session-open-mode [{active-only | both | passive-only}]
```

| Syntax Description | |
|---------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| active-only | Ensures that the BGP session can be established only when the request is initiated by the local end (active-open request) and all passive-open requests (from the other end) are rejected by the local BGP. |
| both | Allows BGP sessions to be established from both incoming or outgoing TCP connection requests, with one being rejected in the event of a request collision. |
| passive-only | Ensures that the local BGP does not initiate any TCP open requests and the session can be established only when the request comes from the remote end. |

Command Default The default is **both** .

Command Modes Neighbor configuration
Neighbor group configuration
Session group configuration

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | 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.

BGP, by default, tries to initiate an active TCP connection whenever a new neighbor is configured. A remote neighbor may also initiate the TCP connection before the local BGP can initiate the connection. This initiation of a TCP connection by a remote neighbor is considered a passive-open request and it is accepted by the local BGP. This default behavior can be modified using the **session-open-mode** command.



Note The BGP connection is not opened and, as a result the BGP session, is not established if both the peering neighbors use the same nondefault TCP session open mode—active-only or passive-only. If both ends are configured with active-only, each neighbor rejects the TCP open request from the other end. One neighbor must be configured as passive-only or both. Similarly, if both neighbors are configured with passive-only, neither neighbor initiates the TCP open request and the BGP session is not established. Again, one neighbor must be configured as active-only or both. There is one exception. A connection open request from a neighbor that is configured with the TCP session open mode to be passive-only is processed to detect whether there is a connection collision before the request is rejected. This exception enables the local BGP to reset the session if the remote neighbor goes down and it is not detected by the local router.

Use the **session-open-mode** command when it may be necessary to preconfigure a neighbor that does not exist. Ensure that BGP does not spend any time actively trying to set up a TCP session with the neighbor. A BGP session does not come up between two neighbors, both of which configure the same nondefault value (**active-only** or **passive-only** keyword) for this command.

Task ID

| Task ID | Operations |
|---------|----------------|
| bgp | read, write |

Examples

The following example shows how to enable a BGP session on router bgp 1:

```
RP/0/RP0/CPU0:router(config)# router bgp 1
RP/0/RP0/CPU0:router(config-bgp)# neighbor 45.67.89.01
RP/0/RP0/CPU0:router(config-bgp-nbr)# session-open-mode active-only
```

show bgp

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

```
show bgp [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. |
| 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. |
| <i>ip-address</i> | (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. |
| <i>mask</i> | (Optional) Network mask of the BGP route to match. |
| <i>/prefix-length</i> | (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 <i>ip-address</i> and <i>mask</i> or <i>/prefix-length</i> arguments are specified. |
| unknown-attributes | (Optional) Includes unknown, transitive attributes. The unknown-attributes keyword is available when the <i>ip-address</i> and <i>mask</i> or <i>/prefix-length</i> arguments are specified. |
| bestpath-compare | (Optional) Displays route and best-path comparison information. The bestpath-compare keyword is available when the <i>ip-address</i> and <i>mask</i> or <i>/prefix-length</i> arguments are specified. |
| standby | (Optional) Displays information about the standby card. |
| detail | (Optional) Displays the prefix details. |

Command Default If no address family or subaddress family is specified, the default address family and subaddress family specified using the **set default-afi** and **set default-safi** commands are used.

Command Modes XR EXEC mode

| Command History | Release | Modification |
|-----------------|---------------|------------------------------------------------------------------------------------------|
| | Release 5.0.0 | This command was introduced. |
| | Release 5.2.1 | The command output was modified to include information from rt-set configuration. |

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 NCS 6000 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 Operations ID

| | |
|-----|------|
| 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 XR EXEC mode:

```

RP/0/RP0/CPU0:router#show bgp
BGP router identifier 172.20.1.1, local AS number 1820
BGP generic scan interval 60 secs
BGP table state: Active
Table ID: 0xe0000000
BGP main routing table version 3
Dampening enabled
BGP scan interval 60 secs

Status codes: s suppressed, d damped, h history, * valid, > best
               i - internal, S stale
Origin codes: i - IGP, e - EGP, ? - incomplete

Network          Next Hop          Metric LocPrf Weight Path
* i10.3.0.0/16   172.20.22.1      0      100      0 1800 1239 ?
*>i              172.20.16.1      0      100      0 1800 1239 ?
* i10.6.0.0/16   172.20.22.1      0      100      0 1800 690 568 ?
*>i              172.20.16.1      0      100      0 1800 690 568 ?
* i10.7.0.0/16   172.20.22.1      0      100      0 1800 701 35 ?
*>i              172.20.16.1      0      100      0 1800 701 35 ?
*                  192.168.40.24    0      100      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      100      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      100      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      100      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. <ul style="list-style-type: none"> • Range for 2-byte Autonomous system numbers (ASNs) is 1 to 65535. • Range for 4-byte Autonomous system numbers (ASNs) in asplain format is 1 to 4294967295. • Range for 4-byte Autonomous system numbers (ASNs) in asdot format is 1.0 to 65535.65535. |

| Field | Description |
|--------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| BGP generic scan interval | Interval (in seconds) between scans of the BGP table by a generic scanner. |
| BGP table state | State of the BGP database. |
| Table ID | BGP database identifier. |
| BGP main routing table version | Last version of the BGP database that was installed into the main routing table. |
| Dampening enabled | Dampening is enabled for the routes in this BGP routing table. |
| BGP scan interval | Interval (in seconds) between BGP scans for the specified address family and subaddress family. |
| Status codes | <p>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):</p> <p>S—Path is stale, indicating that a graceful restart is in progress with the peer from which the route was learned.</p> <p>s—Path is more specific than a locally sourced aggregate route and has been suppressed.</p> <p>*—Path is valid.</p> <p>The second character may be (in order of precedence):</p> <p>>—Path is the best path to use for that network.</p> <p>d—Path is dampened.</p> <p>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.</p> <p>The third character may be:</p> <p>i—Path was learned by an internal BGP (iBGP) session.</p> |
| Origin codes | <p>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:</p> <p>i—Path originated from an Interior Gateway Protocol (IGP) and was sourced by BGP using a network or aggregate-address command.</p> <p>e—Path originated from an Exterior Gateway Protocol (EGP).</p> <p>?—Origin of the path is not clear. Usually, this is a route that is redistributed into BGP from an IGP.</p> |
| Network | IP prefix and prefix length for a network. |
| Next Hop | IP address of the next system that is used when a packet is forwarded to the destination network. An entry of 0.0.0.0 indicates that the router has a non-BGP route to this network. |

| Field | Description |
|--------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Metric | Value of the interautonomous system metric, otherwise known as the Multi Exit discriminator (MED) metric. |
| LocPrf | Local preference value. This is used to determine the preferred exit point from the local autonomous system. It is propagated throughout the local autonomous system. |
| Weight | Path weight. Weight is used in choosing the preferred path to a route. It is not advertised to any neighbor. |
| Path | Autonomous system path to the destination network. At the end of the path is the path origin code. |

The following is sample output from the **show bgp** command with the network specified:

```
RP/0/RP0/CPU0:router# show bgp 11.0.0.0/24
BGP router table entry for 11.0.0.0/24
Versions:
  Process          bRIB/RIB  SendTblVer
  Speaker          2         2

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. |

| Field | Description |
|-----------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 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. |
| 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. |

| Field | Description |
|--------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <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. |
| 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. |

| Field | Description |
|-----------------------|-------------------------------------------------------------------------|
| 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/RP0/CPU0:router# show bgp 51.0.0.0/24 detail
Sat Mar 14 00:37:14.109 PST PDT
BGP routing table entry for 51.0.0.0/24
Versions:
  Process          bRIB/RIB  SendTblVer
  Speaker          3         3
  Flags: 0x3e1000, label_retention: not enabled
Last Modified: Mar 13 19:32:17.976 for 05:04:56
Paths: (1 available, best #1)
  Advertised to update-groups (with more than one peer):
    0.3 0.4 0.7 0.8
  Advertised to peers (in unique update groups):
    201.48.20.1
Path #1: Received by speaker 0
Flags: 0x1000003
200 201
213.0.0.6 from 213.0.0.6 (200.200.3.1)
  Origin IGP, localpref 100, valid, external, best
```

The following is sample output from the **show bgp** command with the additional paths received from:

```
BGP routing table entry for 51.0.1.0/24, Route Distinguisher: 2:1
Versions:
  Process          bRIB/RIB  SendTblVer
  Speaker          63        63
```

```

Flags: 0x040630f2
Last Modified: Nov 11 12:44:05.811 for 00:00:16
Paths: (3 available, best #2)
  Advertised to CE peers (in unique update groups):
    10.51.0.10
  Path #1: Received by speaker 0
  Flags: 0x3
  Not advertised to any peer
  111 111 111 111 111 111 111 111
    10.51.0.10 from 10.51.0.10 (11.11.11.11)
      Origin IGP, metric 0, localpref 100, valid, external
      Received Path ID 0, Local Path ID 0, version 0
      Extended community: RT:55:1
  Path #2: Received by speaker 0
  Flags: 0x5060007
  Advertised to CE peers (in unique update groups):
    10.51.0.10
  561 562 563 564 565
    13.0.6.50 from 13.0.6.50 (13.0.6.50)
      Received Label 16
      Origin IGP, localpref 100, valid, internal, best, group-best, import-candidate,
imported
      Received Path ID 0, Local Path ID 1, version 63
      Extended community: RT:55:1
  Path #3: Received by speaker 0
  Flags: 0x4060007
  Not advertised to any peer
  591 592 593 594 595
    13.0.9.50 from 13.0.9.50 (13.0.9.50)
      Received Label 16
      Origin IGP, localpref 100, valid, internal, backup, add-path, import-candidate,
imported
      Received Path ID 0, Local Path ID 4, version 63
      Extended community: RT:22:232 RT:55:1

```

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 suspect'.

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

```

RP/0/RP0/CPU0:router#show bgp 3.3.3.3/32
Thu Mar 23 14:36:15.301 PDT
BGP routing table entry for 3.3.3.3/32
Versions:
  Process          bRIB/RIB  SendTblVer
  Speaker          42        42
Last Modified: Mar 22 20:04:46.000 for 18:31:30
Paths: (1 available, best #1)
  Advertised to peers (in unique update groups):
    11.11.11.5
  Path #1: Received by speaker 0
  Advertised to peers (in unique update groups):
    11.11.11.5
  Local
    10.10.10.3 from 10.10.10.3 (3.3.3.3)
      Origin incomplete, metric 0, localpref 94, valid, internal, best, group-best
      Received Path ID 0, Local Path ID 0, version 42
      Community: 258:259 260:261 262:263 264:265
      Large Community: 1:2:3 5:6:7 4123456789:4123456780:4123456788

```

Related Commands

| Command | Description |
|--------------------------------------------------------------|---------------------------------------------------------------------------------------------------------|
| aggregate-address , on page 12 | Creates an aggregate entry in a BGP routing table. |
| bgp default local-preference , on page 54 | Changes the default local preference value. |
| network (BGP) , on page 143 | Specifies a local network that the BGP routing process should originate and advertise to its neighbors. |
| route-policy (BGP) , on page 178 | 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 226 | Displays routes with nonnatural netmasks. |
| show bgp community , on page 230 | Displays routes belonging to the specified communities. |
| show bgp inconsistent-as , on page 247 | Displays networks with inconsistent origin autonomous system. |
| show bgp regexp , on page 308 | Displays routes matching an AS path regular expression. |
| show bgp route-policy , on page 312 | Displays networks that match a route policy. |
| show bgp summary , on page 320 | Displays the status of all BGP connections. |
| show bgp truncated-communities , on page 329 | Displays networks with community lists truncated by policy. |

show bgp large-community

To display routes that have the specified Border Gateway Protocol (BGP) large-communities, use the **show bgp large-community** command in XR EXEC mode.

```
show bgp [ ipv4 { unicast | labeled-unicast | all } ] large-community large-community-list
[ exact-match ]
show bgp [ ipv6 { unicast | labeled-unicast | all } ] large-community large-community-list
[ exact-match ]
show bgp [ all { unicast | labeled-unicast | all } ] large-community large-community-list [
exact-match ]
```

| Syntax Description | |
|-----------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| ipv4 | (Optional) Specifies IP Version 4 address prefixes. |
| unicast | (Optional) Specifies unicast 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. |
| large-community | Specifies that only routes with large-communities specified by <i>large-community-list</i> is displayed. |
| <i>large-community-list</i> | 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 XR EXEC

Usage Guidelines To use this command, you must be in a user group associated with a task group that 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 XR EXEC mode:

```
RP/0/RP0/CPU0:router#show bgp large-community 1:2:3 5:6:7
Thu Mar 23 14:40:33.597 PDT
BGP router identifier 4.4.4.4, local AS number 3
BGP generic scan interval 60 secs
Non-stop routing is enabled
BGP table state: Active
Table ID: 0xe0000000 RD version: 66
BGP main routing table version 66
BGP NSR Initial initsync version 3 (Reached)
BGP NSR/ISSU Sync-Group versions 66/0
BGP scan interval 60 secs

Status codes: s suppressed, d damped, h history, * valid, > best
               i - internal, r RIB-failure, S stale, N Nexthop-discard
Origin codes: i - IGP, e - EGP, ? - incomplete
   Network          Next Hop          Metric LocPrf Weight Path
*>i3.3.3.3/32      10.10.10.3          0      94      0 ?
*> 5.5.5.5/32      11.11.11.5          0              0 5 ?
```

In the above example, the routes 3.3.3.3/32 and 5.5.5.5/32 contain both the large communities 1:2:3 and 5:6:7 and some additional large communities.

```
RP/0/RP0/CPU0:router#show bgp large-community 1:2:3 5:6:7 exact-match
Thu Mar 23 14:50:28.155 PDT
```

In the above example, no routes are displayed because the available routes have some additional large-communities, so the match is not exact.

Related Commands

| Command | Description |
|-------------------------------------------------|------------------------------------------------------------------------------------|
| route-policy (BGP), on page 178 | Applies a routing policy to updates advertised to or received from a BGP neighbor. |
| show bgp, on page 193 | Displays BGP routes. |

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 |
|----------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

| Command History | Release | Modification |
|------------------------|----------------|------------------------------------------|
| | Release 5.0.0 | This command was introduced. |
| | Release 6.3.1 | This command supports large communities. |

| | |
|-------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Usage Guidelines | To use this command, you must be in a user group associated with a task group that 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.

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 |

Related Commands

| Command | Description |
|-----------------------------------------------------------|-------------------------------------------------------------------------------------------------------|
| af-group, on page 10 | Creates an address family group for BGP neighbors and enters address family group configuration mode. |
| neighbor-group, on page 141 | Creates a neighbor group and enters neighbor group configuration mode. |
| send-extended-community-ebgp, on page 187 | Specifies that extended community attributes are sent to eBGP neighbors. |

show bgp update out

To display address-family level update generation information, use the **show bgp update out** command in XR EXEC mode.

```
show bgp [afi safi] update out [{brief | detail}]
```

| Syntax Description | |
|--------------------|------------------------------------------------------------------------------|
| <i>afi</i> | (Optional) Displays address-family identifier. |
| <i>safi</i> | (Optional) Displays subsequent address family identifier. |
| brief | (Optional) Displays brief information on process level update generation. |
| detail | (Optional) Displays detailed information on process level update generation. |

Command Default None

Command Modes XR EXEC

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | 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 |
|---------|---------|-----------|
| | bgp | read |

show bgp update in error process

To display process level update inbound error-handling information, use the **show bgp update in error process** command in XR EXEC mode.

show bgp update in error process [{**brief** | **detail**}]

| Syntax Description | |
|--------------------|------------------------------------------------------------------------------|
| brief | (Optional) Displays brief information on process level update generation. |
| detail | (Optional) Displays detailed information on process level update generation. |

Command Default None

Command Modes XR EXEC

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | 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 | Operations |
|---------|---------|------------|
| | bgp | read |

This example displays sample output from the **show bgp update in error process** command:

```
RP/0/RP0/CPU0:router#show bgp update in error process

Basic Update error-handling:
  EBGP: [Enabled]
  IBGP: [Enabled]
Extended Update error-handling:
  EBGP: [Disabled]
  IBGP: [Disabled]

Malformed Update messages: 0
Neighbors that received malformed Update messages: 0
Last malformed Update received: --- (never)
```

show bgp update out filter-group

To display update generation information at filter-group level, **show bgp update out filter-group** command in XR EXEC mode.

```
show bgp [afi safi] update out filter-group [fg-process-id] [{brief | detail}]
```

| Syntax Description | |
|----------------------|----------------------------------------------------------------------------------------|
| <i>afi safi</i> | Specifies the address family and subsequent address family identifiers. |
| <i>fg-process-id</i> | Specifies the filter-group process ID in <x.y> format. Range is <0-15>.<0-4294967295>. |
| brief | (Optional) Displays brief information on filter-group level update generation |
| detail | (Optional) Displays detailed information on filter-group level update generation. |

Command Default None

Command Modes XR EXEC

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | 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 | Operations |
|---------|---------|------------|
| | bgp | read |

This example displays sample output from **show bgp update out filter-group** command:

show bgp update out process

To display process level update generation information, use the **show bgp update out process** command in XR EXEC mode.

show bgp update out process [{**brief** | **detail**}]

| Syntax Description | |
|--------------------|------------------------------------------------------------------------------|
| brief | (Optional) Displays brief information on process level update generation. |
| detail | (Optional) Displays detailed information on process level update generation. |

Command Default None

Command Modes XR EXEC

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | 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 |
|---------|---------|-----------|
| | bgp | read |

This example displays sample output from the show bgp update out process brief command:

```
RP/0/RP0/CPU0:router#show bgp update out process
Wed Sep 12 08:26:04.308 DST

Update generation status: Normal
Update OutQ:          0 bytes (0 messages)
Update limit: 536870912 bytes (configured 536870912 bytes)

Update generation logging: [Disabled]

  Address-family Status   Limit      OutQ      UG   SG (Thr)   SG-R (Thr)  Nbrs
-----
IPv4 Unicast   Normal    268435456  0           1    1(0)       0(0)        3

History:
Update OutQ Hi:          300 bytes (1 messages)
Update OutQ Cumulative: 1200 bytes (4 messages)
Update OutQ Discarded:   0 bytes (0 messages)
Update OutQ Cleared:     0 bytes (0 messages)
Last discarded from OutQ: --- (never)
Last cleared from OutQ:  --- (never)
Update generation throttled 0 times, last event --- (never)
```

```
Update generation recovered 0 times, last event --- (never)
Update generation mem alloc failed 0 times, last event --- (never)
```

show bgp update out sub-group

To display sub-group update generation information, use the **show bgp update out sub-group** command in XR EXEC mode.

```
show bgp [afi safi] update out [update-group ug-index] sub-group [sg-index] [{brief | detail}]
```

Syntax Description

| | |
|-----------------|------------------------------------------------------------------------------|
| afi | (Optional) Displays address-family identifier. |
| safi | (Optional) Displays subsequent address family identifier. |
| brief | (Optional) Displays brief information on process level update generation. |
| detail | (Optional) Displays detailed information on process level update generation. |
| ug-index | (Optional) Displays the update-group process ID in <x.y> format. |
| sg-index | (Optional) displays the sub-group process ID in <x.y> format. |

Command Default

None

Command Modes

XR EXEC

Command History

| Release | Modification |
|---------------|------------------------------|
| 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 |
|---------|-----------|
| bgp | read |

show bgp update out update-group

To display update-group update generation information, use the **show bgp update out update-group** command in XR EXEC mode.

```
show bgp [afi safi] update out update-group [ug-index] [{brief | detail}]
```

| Syntax Description | |
|--------------------|------------------------------------------------------------------------------|
| <i>afi</i> | (Optional) Displays address-family identifier. |
| <i>safi</i> | (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. |
| <i>ug-index</i> | (Optional) Displays the update-group process ID in <x.y> format. |

| | |
|------------------------|------|
| Command Default | None |
|------------------------|------|

| | |
|----------------------|---------|
| Command Modes | XR EXEC |
|----------------------|---------|

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | 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 |
|---------|---------|-----------|
| | bgp | read |

show bgp advertised

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

```
show bgp [ipv4 { all | labeled-unicast | unicast }] advertised [neighbor ip-address] [standby]
[summary]
```

```
show bgp [ipv6 { all | labeled-unicast | unicast}] advertised [neighbor ip-address] [standby]
[summary]
```

```
show bgp [all { all | labeled-unicast | unicast }] advertised [neighbor ip-address] [standby]
[summary]
```

Syntax Description

| | |
|------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------|
| ipv4 | (Optional) Specifies IP Version 4 address prefixes. |
| unicast | (Optional) Specifies unicast address prefixes. |
| labeled-unicast | (Optional) Specifies labeled unicast address prefixes. |
| all | (Optional) For address family, specifies prefixes for all address families. |
| ipv6 | (Optional) Specifies IP Version 6 address prefixes. |
| neighbor | (Optional) Previews advertisements for a single neighbor. If the neighbor keyword is omitted, then the advertisements for all neighbors are displayed. |
| <i>ip-address</i> | (Optional) IP address of the neighbor. |
| standby | (Optional) Displays information about the standby card. |
| summary | (Optional) Displays a summary of advertisements. |

Command Default

If no address family or subaddress family is specified, the default address family and subaddress family specified using the **set default-afi** and **set default-safi** commands are used.

Command Modes

XR EXEC

Command History

| Release | Modification |
|---------------|------------------------------|
| 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.



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 NCS 6000 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.

Task ID

| Task ID | Operations |
|---------|------------|
|---------|------------|

| | |
|-----|------|
| bgp | read |
|-----|------|

Examples

The following is sample output from the **show bgp advertised** command in XR EXEC mode:

```
RP/0/RP0/CPU0:router# show bgp advertised neighbor 10.0.101.4 summary
```

| Network | Next Hop | From | AS Path |
|------------|------------|------------|---------------------------|
| 1.1.1.0/24 | 10.0.101.1 | 10.0.101.1 | 2 3 222 333 444 555 i |
| 1.1.2.0/24 | 10.0.101.1 | 10.0.101.1 | 3 4 5 6 7 i |
| 1.1.3.0/24 | 10.0.101.1 | 10.0.101.1 | 77 88 33 44 55 99 99 99 i |
| 1.1.4.0/24 | 10.0.101.1 | 10.0.101.1 | 2 5 6 7 8 i |
| 1.1.7.0/24 | 10.0.101.1 | 10.0.101.1 | 3 5 i |
| 1.1.8.0/24 | 10.0.101.1 | 10.0.101.1 | 77 88 99 99 99 i |

This table describes the significant fields shown in the display.

Table 5: show bgp advertised neighbor summary Field Descriptions

| Field | Description |
|-----------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Network | IP prefix and prefix length for a network. |
| Next Hop | IP address of the next system that is used when a packet is forwarded to the destination network. An entry of 0.0.0.0 indicates that the router has a non-BGP route to this network. |
| From | IP address of the peer that advertised this route. |
| AS Path | AS path of the peer that advertised this route. |
| Local | Indicates the route originated on the local system. |
| Local Aggregate | Indicates the route is an aggregate created on the local system. |
| Advertised to | Indicates the peer to which this entry was advertised. This field is used in the output when displaying a summary of the advertisements to all neighbors. |

The following is sample output from the **show bgp advertised** command for detailed advertisement information:

```
RP/0/RP0/CPU0:router# show bgp advertised neighbor 172.72.77.1

172.16.0.0/24 is advertised to 172.72.77.1
  Path info:
    neighbor: Local          neighbor router id: 172.74.84.1
    valid redistributed best
  Attributes after inbound policy was applied:
  next hop: 0.0.0.0
    MET ORG AS
    origin: incomplete metric: 0
    aspath:
10.52.0.0/16 is advertised to 172.72.77.1
  Path info:
    neighbor: Local Aggregate neighbor router id: 172.74.84.1
    valid aggregated best
  Attributes after inbound policy was applied:
  next hop: 0.0.0.0
    ORG AGG ATOM
    origin: IGP aggregator: 172.74.84.1 (1)
    aspath:
```

This table describes the significant fields shown in the display.

Table 6: show bgp advertised neighbor Field Descriptions

| Field | Description |
|------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| is advertised to | IP address of the peer to which this route has been advertised. If the route has been advertised to multiple peers, the information is shown separately for each peer. |

| 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 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. |
| (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. |

| Field | Description |
|---------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| dampinfo | <p>Indicates dampening information:</p> <p>Penalty—Current penalty for this path.</p> <p>Flapped—Number of times the route has flapped.</p> <p>In—Time (hours:minutes:seconds) since the router noticed the first flap.</p> <p>Reuse in—Time (hours:minutes:seconds) after which the path is made available. This field is displayed only if the path is currently suppressed.</p> |
| Attributes after inbound policy was applied | <p>Displays attributes associated with the received route, after any inbound policy has been applied.</p> <p>AGG—Aggregator attribute is present.</p> <p>AS—AS path attribute is present.</p> <p>ATOM—Atomic aggregate attribute is present.</p> <p>COMM—Communities attribute is present.</p> <p>EXTCOMM—Extended communities attribute is present.</p> <p>LOCAL—Local preference attribute is present.</p> <p>MET—Multi Exit Discriminator (MED) attribute is present.</p> <p>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.</p> <p>ORG—Origin attribute is present.</p> |
| origin | <p>Origin of the path:</p> <p>IGP—Path originated from an Interior Gateway Protocol (IGP) and was sourced by BGP using a network or aggregate-address command.</p> <p>EGP—Path originated from an Exterior Gateway Protocol.</p> <p>incomplete—Origin of the path is not clear. For example, a route that is redistributed into BGP from an IGP.</p> |
| 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 | <p>Community attributes associated with the path. Community values are displayed in AA:NN format, except for the following well-known communities:</p> <p>Local-AS—Community with value 4294967043 or hex 0xFFFFFFFF03. Routes with this community value are not advertised outside the local autonomous system or confederation boundary.</p> <p>no-advertise—Community with value 4294967042 or hex 0xFFFFFFFF02. Routes with this community value are not advertised to any BGP peers.</p> <p>no-export—Community with value 4294967041 or hex 0xFFFFFFFF01. Routes with this community are not advertised to external BGP peers, even if those peers are in the same confederation with the local router.</p> |
| Extended community | <p>Extended community attributes associated with the path. For known extended community types, the following codes may be displayed:</p> <p>RT—Route target community</p> <p>SoO—Site of Origin community</p> <p>LB—Link Bandwidth community</p> |
| 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 178 | Applies a route policy to incoming and outgoing routes. |
| show bgp policy, on page 283 | Displays information about BGP advertisements under a proposed policy. |
| sent-advertisements | Previews advertisements to peers, including details of advertised attributes. |

show bgp af-group

To display information about Border Gateway Protocol (BGP) configuration for address family groups, use the **show bgp af-group** command in XR EXEC mode.

```
show bgp af-group group-name {configuration [defaults] [nvgen] | inheritance | users}
```

Syntax Description

| | |
|----------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>group-name</i> | Name of the address family group to display. |
| configuration | (Optional) Displays the effective configuration for the af-group, including any settings that have been inherited from af-groups used by this af-group. |
| defaults | (Optional) Displays all configuration settings, including any default settings. |
| nvgen | (Optional) Displays output in the format of show running-config output. If the defaults keyword is also specified, the output is not suitable for cutting and pasting into a configuration session. |
| inheritance | Displays the af-groups from which this af-group inherits configuration settings. |
| users | Displays the neighbors, neighbor groups, and af-groups that inherit configuration from this af-group. |

Command Default

No default behavior or value

Command Modes

XR EXEC

Command History

| Release | Modification |
|---------------|------------------------------|
| 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.

Use the **show bgp af-group** command with the *group-name* **configuration** argument and keyword to display the effective configuration of an af-group, taking into account any configuration that may be inherited from other af-groups through the **use af-group** command. The source of each command is shown.

If the **defaults** keyword is specified, all configuration for the af-group, including default values, is shown. Default configuration is identified in the show output. Use the **nvgen** keyword to display configuration formatted in the style of the **show running-config** command. This output is suitable for cutting and pasting into configuration sessions.

Use the **show bgp af-group** command with the *group-name* **inheritance** argument and keyword to display the address family groups from which the specified af-group inherits configuration.

Use the **show bgp af-group** command with the *group-name* **users** argument and keyword to display the neighbors, neighbor groups, and af-groups that inherit configuration from the specified af-group.

| Task ID | Task ID | Operations |
|---------|---------|------------|
| | bgp | read |

Examples

The following af-group configuration is used in the examples:

```
af-group group3 address-family ipv4 unicast
remove-private-AS
soft-reconfiguration inbound
!
af-group group1 address-family ipv4 unicast
use af-group group2
maximum-prefix 2500 75 warning-only
default-originate
soft-reconfiguration inbound disable
!
af-group group2 address-family ipv4 unicast
use af-group group3
send-community-ebgp
send-extended-community-ebgp
capability orf prefix both
```

The following is sample output from the **show bgp af-group** command with the **configuration** keyword in XR EXEC mode. The source of each command is shown in the right column. For example, **default-originate** is configured directly on **af-group group1**, and the **remove-private-AS** command is inherited from af-group group2, which in turn inherits it from af-group group3.

```
RP/0/RP0/CPU0:router# show bgp af-group group1 configuration

af-group group1 address-family ipv4 unicast
  capability orf prefix both                [a:group2]
  default-originate                         []
  maximum-prefix 2500 75 warning-only      []
  remove-private-AS                        [a:group2 a:group3]
  send-community                            [a:group2]
  send-extended-community                   [a:group2]
```

The following is sample output from the **show bgp af-group** command with the **users** keyword:

```
RP/0/RP0/CPU0:router# show bgp af-group group2 users

IPv4 Unicast: a:group1
```

The following is sample output from the **show bgp af-group** command with the **inheritance** keyword. This example shows that the specified af-group group1 directly uses the group2 af-group, which in turn uses the group3 af-group:

```
RP/0/RSP0RP0/CPU0:router# show bgp af-group group1 inheritance

IPv4 Unicast: a:group2 a:group3
```

Table 7: [show bgp af-group Field Descriptions, on page 222](#) describes the significant fields shown in the display.

This table describes the significant fields shown in the display.

Table 7: show bgp af-group Field Descriptions

| Field | Description |
|-----------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| [] | Configures the command directly on the specified address family group. |
| a: | Indicates the name that follows is an address family group. |
| n: | Indicates the name that follows is a neighbor group. |
| [dflt] | Indicates the setting is not explicitly configured or inherited, and the default value for the setting is used. This field may be shown when the defaults keyword is specified. |
| <not set> | Indicates that the configuration is disabled by default. This field may be shown when the defaults keyword is specified. |

Related Commands

| Command | Description |
|------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------|
| af-group, on page 10 | Configures a BGP address family group. |
| show bgp neighbors, on page 256 | Displays information about BGP neighbors, including configuration inherited from neighbor groups, session groups, and address family groups. |
| show bgp neighbor-group, on page 252 | Displays information about configuration for neighbor groups. |
| use, on page 368 af-group | Configures an af-group to inherit the configuration of a specified af-group. |

show bgp attribute-key

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

```
show bgp {ipv4 | ipv6 | all} attribute-key [standby]
```

| Syntax Description | |
|------------------------|-----------------------------------------------------------------------------------|
| ipv4 | (Optional) Specifies IP Version 4 address prefixes. |
| unicast | (Optional) Specifies unicast address prefixes. |
| labeled-unicast | (Optional) Specifies labeled unicast address prefixes. |
| all | (Optional) For address family, specifies prefixes for all address families. |
| all | (Optional) For subaddress family, specifies prefixes for all subaddress families. |
| ipv6 | (Optional) Specifies IP Version 6 address prefixes. |
| standby | (Optional) Displays information about the standby card. |

Command Default If no address family or subaddress family is specified, the default address family and subaddress family specified using the **set default-afi** and **set default-safi** commands are used.

Command Modes XR EXEC

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | 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.



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 NCS 6000 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 XR EXEC mode:

show bgp attribute-key

```

RP/0/RP0/CPU0:router# show bgp all all attribute-key

Address Family: IPv4 Unicast
=====

BGP router identifier 10.0.0.1, local AS number 1
BGP generic scan interval 60 secs
BGP main routing table version 109
BGP scan interval 60 secs
Status codes: s suppressed, d damped, h history, * valid, > best
               i - internal, S stale
Origin codes: i - IGP, e - EGP, ? - incomplete
   Network      Next Hop      AttrKey
*> 1.1.0.0/16   0.0.0.0       0x00000002
*> 10.0.0.0/16  0.0.0.0       0x00000002
*> 12.21.0.0/16 0.0.0.0       0x00000002
*> 194.3.192.1/32 10.0.101.1    0x00000009
*> 194.3.192.2/32 10.0.101.1    0x00000009
*> 194.3.192.3/32 10.0.101.1    0x00000009
*> 194.3.192.4/32 10.0.101.1    0x00000009
*> 194.3.192.5/32 10.0.101.1    0x00000009

Processed 8 prefixes, 8 paths

Address Family: IPv6 Unicast
=====

BGP router identifier 10.0.0.1, local AS number 1
BGP generic scan interval 60 secs
BGP main routing table version 19
BGP scan interval 60 secs
Status codes: s suppressed, d damped, h history, * valid, > best
               i - internal, S stale
Origin codes: i - IGP, e - EGP, ? - incomplete
   Network      Next Hop      AttrKey
*> 2222::1111/128 2222::2       0x00000009
*> 2222::1112/128 2222::2       0x00000009

Processed 2 prefixes, 2 paths

```

This table describes the significant fields shown in the display.

Table 8: show bgp attribute-key Field Descriptions

| Field | Description |
|--------------------------------|----------------------------------------------------------------------------------|
| BGP router identifier | BGP identifier for the local system. |
| local AS number | Autonomous system number for the local system. |
| BGP generic scan interval | Interval (in seconds) between scans of the BGP table by a generic scanner. |
| BGP main routing table version | Last version of the BGP database that was installed into the main routing table. |
| BGP scan interval | Interval (in seconds) between scans. |

| Field | Description |
|---------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Status codes | <p>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):</p> <p>S—Path is stale, indicating that a graceful restart is in progress with the peer from which the route was learned.</p> <p>s—Path is more specific than a locally sourced aggregate route and has been suppressed.</p> <p>*—Path is valid.</p> <p>The second character may be (in order of precedence):</p> <p>>—Path is the best path to use for that network.</p> <p>d—Path is dampened.</p> <p>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.</p> <p>The third character may be:</p> <p>i—Path was learned by an internal BGP (iBGP) session.</p> |
| Origin codes | <p>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:</p> <p>i—Path originated from an Interior Gateway Protocol (IGP) and was sourced by BGP using a network or aggregate-address command.</p> <p>e—Entry originated from an Exterior Gateway Protocol (EGP).</p> <p>?—Origin of the path is not clear. Usually, this is a route that is redistributed into BGP from an IGP.</p> |
| 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 <i>n</i> prefixes, <i>n</i> 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 cidr-only

To display routes with nonnatural network masks, also known as classless interdomain routing (CIDR) routes, use the **show bgp cidr-only** command in XR EXEC mode.

```
show bgp [{ipv4}] cidr-only [standby]
```

| Syntax Description | | |
|------------------------|------------|-----------------------------------------------------------|
| ipv4 | (Optional) | Specifies the IP Version 4 address family. |
| unicast | (Optional) | Specifies the unicast address family. |
| labeled-unicast | (Optional) | Specifies labeled unicast address prefixes. |
| all | (Optional) | For subaddress family, specifies all subaddress families. |
| standby | (Optional) | Displays information about the standby card. |

Command Default If no address family or subaddress family is specified, the default address family and subaddress family specified using the **set default-afi** and **set default-safi** commands are used. This command is applicable only for IPv4 prefixes. If the default address family is not IPv4, then the **ipv4** keyword must be used.

Command Modes XR EXEC

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | 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.



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 NCS 6000 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 XR EXEC mode:

```
RP/0/RP0/CPU0:router# show bgp cidr-only

BGP router identifier 172.20.1.1, local AS number 1820
BGP main routing table version 2589
Dampening enabled
BGP scan interval 60 secs
Status codes: s suppressed, d damped, h history, * valid, > best
                i - internal, S stale
Origin codes: i - IGP, e - EGP, ? - incomplete
Network        Next Hop           Metric   LocPrf   Weight   Path
*> 192.0.0.0/8  192.168.72.24       0        1878    ?
*> 192.168.0.0/16 192.168.72.30       0        108     ?
```

This table describes the significant fields shown in the display.

Table 9: show bgp cidr-only Field Descriptions

| Field | Description |
|--------------------------------|-------------------------------------------------------------------------------------------------------------|
| BGP router identifier | BGP identifier for the local system. |
| local AS number | Autonomous system number for the local system. |
| BGP main routing table version | Last version of the BGP database that was installed into the main routing table. |
| Dampening enabled | Displayed if dampening is enabled for the routes in this BGP routing table. |
| BGP scan interval | Interval (in seconds) between scans of the BGP table specified by the address family and subaddress family. |

| Field | Description |
|--------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Status codes | <p>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):</p> <p>S—Path is stale, indicating that a graceful restart is in progress with the peer from which the route was learned.</p> <p>s—Path is more specific than a locally sourced aggregate route and has been suppressed.</p> <p>*—Path is valid.</p> <p>The second character may be (in order of precedence):</p> <p>>—Path is the best path to use for that network.</p> <p>d—Path is dampened.</p> <p>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.</p> <p>The third character may be:</p> <p>i—Path was learned by an internal BGP (iBGP) session.</p> |
| Origin codes | <p>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:</p> <p>i—Path originated from an Interior Gateway Protocol (IGP) and was sourced by BGP using a network or aggregate-address command.</p> <p>e—Entry originated from an Exterior Gateway Protocol (EGP).</p> <p>?—Origin of the path is not clear. Usually, this is a route that is redistributed into BGP from an IGP.</p> |
| 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 178 | 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 193 | 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 XR EXEC mode.

```
show bgp [ipv4 {unicast | labeled-unicast | all}] community community-list [exact-match]
show bgp [ipv6 {unicast | labeled-unicast | all}] community community-list [exact-match]
show bgp [all {unicast | labeled-unicast | all}] community community-list [exact-match]
```

| Syntax | Description |
|------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| ipv4 | (Optional) Specifies IP Version 4 address prefixes. |
| unicast | (Optional) Specifies unicast 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. |
| community | Specifies that only routes with communities specified by <i>community-list</i> is displayed. |
| <i>community-list</i> | <p>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:</p> <p>local-AS —Well-known community with value 4294967043 or hex 0xFFFFFFFF03. Routes with this community value are not advertised outside the local autonomous system or confederation boundary.</p> <p>no-advertise —Well-known community with value 4294967042 or hex 0xFFFFFFFF02. Routes with this community value are not advertised to any BGP peers.</p> <p>no-export —Well-known community with value 4294967041 or hex 0xFFFFFFFF01. Routes with this community are not advertised to external BGP peers, even if those peers are in the same confederation as the local router.</p> <p>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.</p> <p>For the AA:NN format:</p> <p>AA—Range is 0 to 65535.</p> <p>NN—Range is 1 to 4294967295.</p> <p>Up to seven community numbers can be specified.</p> |
| exact-match | (Optional) Displays those routes that have communities exactly matching the specified communities. |

Command Default If no address family or subaddress family is specified, the default address family and subaddress family specified using the **set default-afi** and **set default-safi** commands are used.

Command Modes XR EXEC

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | 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.



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 NCS 6000 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 312](#) 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
```

show bgp community

```
* 5.5.5.5/32          10.10.10.1          88          0 1 ?
Processed 1 prefixes, 1 paths
```

Examples

The following is sample output from the **show bgp community** command in XR EXEC mode:

```
RP/0/RP0/CPU0:router# show bgp community 1820:1 exact-match

BGP router identifier 172.20.1.1, local AS number 1820
BGP main routing table version 55
Dampening enabled
BGP scan interval 60 secs
Status codes: s suppressed, d damped, h history, * valid, > best
               i - internal, S stale
Origin codes: i - IGP, e - EGP, ? - incomplete
   Network          Next Hop          Metric LocPrf Weight Path
*  10.13.0.0/16     192.168.40.24          0  1878  704  701  200  ?
*  10.16.0.0/16     192.168.40.24          0  1878  704  701  i
```

This table describes the significant fields shown in the display.

Table 10: show bgp community Field Descriptions

| Field | Description |
|--------------------------------|-------------------------------------------------------------------------------------------------------------|
| BGP router identifier | BGP identifier for the local system. |
| local AS number | Autonomous system number for the local system. |
| 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 | <p>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):</p> <p>S—Path is stale, indicating that a graceful restart is in progress with the peer from which the route was learned.</p> <p>s—Path is more specific than a locally sourced aggregate route and has been suppressed.</p> <p>*—Path is valid.</p> <p>The second character may be (in order of precedence):</p> <p>>—Path is the best path to use for that network.</p> <p>d—Path is dampened.</p> <p>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.</p> <p>The third character may be:</p> <p>i—Path was learned by an internal BGP (iBGP) session.</p> |
| Origin codes | <p>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:</p> <p>i—Path originated from an Interior Gateway Protocol (IGP) and was advertised with a network or aggregate-address command.</p> <p>e—Path originated from an Exterior Gateway Protocol (EGP).</p> <p>?—Origin of the path is not clear. Usually, this is a route that is redistributed into BGP from an IGP.</p> |
| 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 12 | Creates an aggregate entry in a BGP routing table. |
| network (BGP), on page 143 | Specifies a local network that the BGP routing process should originate and advertise to its neighbors. |
| route-policy (BGP), on page 178 | 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 193 | Displays BGP routes. |

show bgp convergence

To display whether a specific address family has reached convergence, use the **show bgp convergence** command in XR EXEC mode.

```
show bgp [ipv4 {unicast | labeled-unicast | all}] convergence
show bgp [ipv6 {unicast | labeled-unicast | all}] convergence
show bgp [all {unicast | labeled-unicast | all}] convergence
```

| Syntax Description | ipv4 (Optional) Specifies the IP Version 4 address family. | | | | |
|-------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------|--------------|---------------|------------------------------|
| | unicast (Optional) Specifies the unicast address family. | | | | |
| | labeled-unicast (Optional) Specifies unicast address prefixes. | | | | |
| | all (Optional) For subaddress family, specifies all subaddress families. | | | | |
| | ipv6 (Optional) Specifies the IP Version 6 address family. | | | | |
| | all (Optional) For address family, specifies all address families. | | | | |
| Command Default | If no address family or subaddress family is specified, the default address family and subaddress family specified using the set default-afi and set default-safi commands are used. | | | | |
| Command Modes | XR EXEC | | | | |
| Command History | <table border="1"> <thead> <tr> <th>Release</th> <th>Modification</th> </tr> </thead> <tbody> <tr> <td>Release 5.0.0</td> <td>This command was introduced.</td> </tr> </tbody> </table> | Release | Modification | Release 5.0.0 | This command was introduced. |
| Release | Modification | | | | |
| 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. | | | | |
|  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 <i>System Management Command Reference for Cisco NCS 6000 Series Routers</i> for detailed information and syntax for the set default-afi and set default-safi commands. If you do not specify a default address family, the default address family is IPv4. If you do not specify a default subaddress family, the default subaddress family is unicast. | | | | |
| | Border Gateway Protocol (BGP) contains a separate routing table for each configured address family and subaddress family combination. The address family and subaddress family options specify the routing table to be examined. If the all keyword is specified for the address family or subaddress family, each matching routing table is examined in turn. | | | | |

Use the **show bgp convergence** command to see if there is any pending work for BGP to perform. The software checks the following conditions to determine whether the specified address family has converged. If all the conditions are true, the address family is considered converged.

- All received updates have been processed and best routes selected.
- All selected routes have been installed in the global Routing Information Base (RIB).
- All selected routes have been advertised to peers, including any peers that are not established (unless those peers have been administratively shut down). See the **shutdown (BGP)** command for more information about administrative shutdown.

While testing that all selected routes have been advertised to peers, the **show bgp convergence** command checks the size of the write queue for each neighbor. Because this queue is shared by all address families, there is a small possibility that the command indicates the address family has not converged when, in fact, it has converged. This could happen if the neighbor write queue contained messages from some other address family.

If the specified address family has not converged, the **show bgp convergence** command output does not indicate the amount of work that is pending. To display this information, use the **show bgp summary** command.

| Task ID | Task ID | Operations |
|---------|---------|------------|
| | bgp | read |

Examples

The following shows the result of using the **show bgp convergence** command for an address family that has converged:

```
RP/0/RP0/CPU0:router# show bgp convergence
Converged.
All received routes in RIB, all neighbors updated.
All neighbors have empty write queues.
```

The following shows the result of using the **show bgp convergence** command for an address family that has not converged:

```
RP/0/RP0/CPU0:router# show bgp convergence
Not converged.
Received routes may not be entered in RIB.
One or more neighbors may need updating.
```

This table describes the significant fields shown in the display.

Table 11: show bgp convergence Field Descriptions

| Field | Description |
|-------------------------|-------------------------------------------------------------------------------------------------------------------------------|
| Converged/Not converged | Specifies whether or not all routes have been installed in the RIB and updates have been generated and sent to all neighbors. |

| Field | Description |
|----------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| [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 320 | Displays the status of BGP peer connections. |
| shutdown (BGP), on page 339 | 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 XR EXEC mode.

```
show bgp [ipv4 {unicast | labeled-unicast | all}] dampened-paths [standby]
show bgp [ipv6 {unicast | labeled-unicast | all}] dampened-paths [standby]
```

| Syntax Description | |
|------------------------|-------------------------------------------------------------------------------------|
| ipv4 | (Optional) Specifies IP Version 4 address prefixes. |
| unicast | (Optional) Specifies unicast 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. |
| standby | (Optional) Displays information about the standby card. |

Command Default If no address family or subaddress family is specified, the default address family and subaddress family specified using the **set default-afi** and **set default-safi** commands are used.

Command Modes XR EXEC

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | 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.



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 NCS 6000 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 XR EXEC mode:

```
RP/0/RP0/CPU0:router# show bgp dampened-paths

BGP router identifier 10.2.0.1, local AS number 3
BGP generic scan interval 60 secs
BGP main routing table version 7
Dampening enabled
BGP scan interval 60 secs
Status codes:s suppressed, d damped, h history, * valid, > best
                i - internal, S stale

Origin codes:i - IGP, e - EGP, ? - incomplete
  Network      From           Reuse      Path
*d 10.0.0.0    10.0.101.35    00:01:20  35 i
```

This table describes the significant fields shown in the display.

Table 12: show bgp dampened-paths Field Descriptions

| Field | Description |
|--------------------------------|-------------------------------------------------------------------------------------------------------------|
| BGP router identifier | BGP identifier for the local system. |
| local AS number | Autonomous system number for the local system. |
| BGP generic scan interval | Interval (in seconds) between scans of the BGP table by a generic scanner. |
| BGP main routing table version | Last version of the BGP database that was installed into the main routing table. |
| Dampening enabled | Displayed if dampening is enabled for the routes in this BGP routing table. |
| BGP scan interval | Interval (in seconds) between scans of the BGP table specified by the address family and subaddress family. |

| Field | Description |
|--------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Status codes | <p>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):</p> <p>S—Path is stale, indicating that a graceful restart is in progress with the peer from which the route was learned.</p> <p>s—Path is more specific than a locally sourced aggregate route and has been suppressed.</p> <p>*—Path is valid.</p> <p>The second character may be (in order of precedence):</p> <p>>—Path is the best path to use for that network.</p> <p>d—Path is dampened.</p> <p>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.</p> <p>The third character may be:</p> <p>i—Path was learned by an internal BGP (iBGP) session.</p> |
| Origin codes | <p>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:</p> <p>i—Path originated from an Interior Gateway Protocol (IGP) and was advertised with a network or aggregate-address command.</p> <p>e—Path originated from an Exterior Gateway Protocol (EGP).</p> <p>?—Origin of the path is not clear. Usually, this is a route that is redistributed into BGP from an IGP.</p> |
| 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 12 | Creates an aggregate entry in a BGP routing table. |
| bgp dampening, on page 52 | Enables BGP route dampening or changes various BGP route dampening factors. |
| clear bgp dampening, on page 85 | Clears BGP route dampening information and unsuppresses the suppressed routes. |

| Command | Description |
|-------------------------------------------------------|---------------------------------------------------------------------------------------------------------|
| network (BGP), on page 143 | 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 242 | Displays BGP routes that have flapped. |
| show bgp neighbors, on page 256 | 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 XR EXEC mode.

```
show bgp [ipv4 {unicast | labeled-unicast | all}] flap-statistics [{regex
regular-expression | route-policy route-policy-name | cidr-only | {ip-address | {mask /prefix-length}}}]
[longer-prefixes] [detail] [standby]
show bgp [ipv6 {unicast | labeled-unicast | all}] flap-statistics [{regex
regular-expression | route-policy route-policy-name | cidr-only | {ip-address | {mask /prefix-length}}}]
[longer-prefixes] [detail] [standby]
show bgp [all {unicast | labeled-unicast | all}] flap-statistics [{regex regular-expression | route-policy
route-policy-name | cidr-only | {ip-address | {mask /prefix-length}}}] [longer-prefixes] [detail]
[standby]
```

Syntax Description

| | |
|-------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| ipv4 | (Optional) Specifies IP Version 4 address prefixes. |
| unicast | (Optional) Specifies unicast 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. |
| regex <i>regular-expression</i> | (Optional) Displays flap statistics for all paths that match the regular expression. |
| route-policy <i>route-policy-name</i> | (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. |
| <i>ip-address</i> | (Optional) Flap statistics for a network address only. |
| <i>mask</i> | (Optional) Network mask applied to the <i>ip-address</i> argument. |
| <i>/ prefix-length</i> | (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 <i>ip-address</i> and <i>mask</i> or <i>/prefix-length</i> 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 <i>ip-address</i> argument or <i>ip-address</i> and <i>mask</i> or <i>/prefix-length</i> arguments are specified. |

standby (Optional) Displays information about the standby card.

Command Default

If no address family or subaddress family is specified, the default address family and subaddress family specified using the **set default-afi** and **set default-safi** commands are used.

Command Modes

XR EXEC

Command History

| Release | Modification |
|---------------|------------------------------|
| 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.

**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 NCS 6000 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 **regex**, **filter-list**, **cidr-only**, and **longer-prefixes** options to limit the set of paths displayed.

If you specify a network address without a mask or prefix length, the longest matching prefix for the network address is displayed. When displaying flap statistics for a single route, use the **detail** keyword to display dampening parameters for the route.

Task ID

| Task ID | Operations |
|---------|------------|
| bgp | read |

Examples

The following is sample output from the **show bgp flap-statistics** command:

```
RP/0/RP0/CPU0:router# show bgp flap-statistics

BGP router identifier 172.20.1.1, local AS number 1820
BGP main routing table version 26180
Dampening enabled
```

```

BGP scan interval 60 secs
Status codes: s suppressed, d damped, h history, * valid, > best
               i - internal, S stale
Origin codes: i - IGP, e - EGP, ? - incomplete
Network      From          Flaps Duration Reuse    Path
*d 10.0.0.0   172.20.16.177  4      00:13:31 00:18:10 100
*d 10.10.0.0  172.20.16.177  4      00:02:45 00:28:20 100

```

The following is sample output from the **show bgp flap-statistics** command with the **detail** keyword in XR EXEC mode:

```

RP/0/RP0/CPU0:router# show bgp flap-statistics 172.31.12.166 detail

BGP router identifier 10.0.0.5, local AS number 1
BGP main routing table version 738
Dampening enabled
BGP scan interval 60 secs
Status codes: s suppressed, d damped, h history, * valid, > best
               i - internal, S stale
Origin codes: i - IGP, e - EGP, ? - incomplete
Network      From          Flaps Duration Reuse    Path
h 172.31.12.166 10.0.101.1    6      00:03:28          2 2000 3000

Half life      Suppress      Reuse penalty  Max. supp. time
00:15:00      2000          750            01:00:00

```

This table describes the significant fields shown in the display.

Table 13: show bgp flap-statistics Field Descriptions

| Field | Description |
|--------------------------------|-------------------------------------------------------------------------------------------------------------|
| BGP route identifier | BGP identifier for the local system. |
| local AS number | Autonomous system number for the local system. |
| BGP main routing table version | Last version of the BGP database that was installed into the main routing table. |
| Dampening enabled | Displayed if dampening has been enabled for the routes in this BGP routing table. |
| BGP scan interval | Interval (in seconds) between scans of the BGP table specified by the address family and subaddress family. |

| Field | Description |
|--------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Status codes | <p>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):</p> <p>S—Path is stale, indicating that a graceful restart is in progress with the peer from which the route was learned.</p> <p>s—Path is more specific than a locally sourced aggregate route and has been suppressed.</p> <p>*—Path is valid.</p> <p>The second character may be (in order of precedence):</p> <p>>—Path is the best path to use for that network.</p> <p>d—Path is dampened.</p> <p>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.</p> <p>The third character may be:</p> <p>i—Path was learned by an internal BGP (iBGP) session.</p> |
| Origin codes | <p>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:</p> <p>i—Path originated from an Interior Gateway Protocol (IGP) and was advertised with a network or aggregate-address command.</p> <p>e—Path originated from an Exterior Gateway Protocol (EGP).</p> <p>?—Origin of the path is not clear. Usually, this is a route that is redistributed into BGP from an IGP.</p> |
| 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. |
| 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. |

| Field | Description |
|----------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 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 52 | 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 238 | Displays the BGP dampened routes. |
| show bgp neighbors, on page 256 | 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 XR EXEC mode.

```
show bgp [ipv4 {unicast | labeled-unicast | all}] inconsistent-as [standby]
show bgp [ipv6 {unicast | labeled-unicast | all}] inconsistent-as [standby]
show bgp [all {unicast | labeled-unicast | all}] inconsistent-as [standby]
```

| Syntax Description | |
|------------------------|-------------------------------------------------------------------------------------|
| ipv4 | (Optional) Specifies IP Version 4 address prefixes. |
| unicast | (Optional) Specifies unicast 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. |
| standby | (Optional) Displays information about the standby card. |

Command Default If no address family or subaddress family is specified, the default address family and subaddress family specified using the **set default-afi** and **set default-safi** commands are used.

Command Modes XR EXEC

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | 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.



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 NCS 6000 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 XR EXEC mode:

```
RP/0/RP0/CPU0:router# show bgp inconsistent-as

BGP router identifier 172.20.1.1, local AS number 1820
BGP main routing table version 1129
BGP scan interval 60 secs
Status codes: s suppressed, d damped, h history, * valid, > best
                i - internal, S stale
Origin codes: i - IGP, e - EGP, ? - incomplete
Network        Next Hop           Metric           LocPrf Weight Path
* 10.0.0.0      172.16.232.55      0                0 300 88 90 99 ?
*>             172.16.232.52      2222             0 400 ?
* 172.16.0.0    172.16.232.55      0                0 300 90 99 88 200 ?
*>             172.16.232.52      2222             0 400 ?
* 192.168.199.0 172.16.232.55      0                0 300 88 90 99 ?
*>             172.16.232.52      2222             0 400 ?
```

This table describes the significant fields shown in the display.

Table 14: show bgp inconsistent-as Field Descriptions

| Field | Description |
|--------------------------------|-------------------------------------------------------------------------------------------------------------|
| BGP router identifier | BGP identifier for the local system. |
| local AS number | Autonomous system number for the local system. |
| BGP main routing table version | Last version of the BGP database that was installed into the main routing table. |
| Dampening enabled | Displayed if dampening is enabled for the routes in this BGP routing table. |
| BGP scan interval | Interval (in seconds) between scans of the BGP table specified by the address family and subaddress family. |

| Field | Description |
|--------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Status codes | <p>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):</p> <p>S—Path is stale, indicating that a graceful restart is in progress with the peer from which the route was learned.</p> <p>s—Path is more specific than a locally sourced aggregate route and has been suppressed.</p> <p>*—Path is valid.</p> <p>The second character may be (in order of precedence):</p> <p>>—Path is the best path to use for that network.</p> <p>d—Path is dampened.</p> <p>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.</p> <p>The third character may be:</p> <p>i—Path was learned by an internal BGP (iBGP) session.</p> |
| Origin codes | <p>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:</p> <p>i—Path originated from an Interior Gateway Protocol (IGP) and was advertised with a network or aggregate-address command.</p> <p>e—Path originated from an Exterior Gateway Protocol (EGP).</p> <p>?—Origin of the path is not clear. Usually, this is a route that is redistributed into BGP from an IGP.</p> |
| 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 |
|-------------------------------------------------|---------------------------------------------------------------------------------------------------------|
| network (BGP), on page 143 | Specifies a local network that the BGP routing process should originate and advertise to its neighbors. |
| route-policy (BGP), on page 178 | 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 labels

To display Border Gateway Protocol (BGP) routes and their incoming and outgoing labels, use the **show bgp labels** command in XR EXEC mode.

show bgp labels

| Syntax | Description |
|------------------------|-------------------------------------------------------------------------------------|
| ipv4 | (Optional) Specifies IP Version 4 address prefixes. |
| unicast | (Optional) Specifies unicast address prefixes. |
| 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. |
| standby | (Optional) Displays information about the standby card. |

Command Default If no address family or subaddress family is specified, the default address family and subaddress family specified using the **set default-afi** and **set default-safi** commands are used.

Command Modes XR EXEC

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | 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 | Operations |
|---------|---------|------------|
| | bgp | read |

| 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 neighbor-group

To display information about the Border Gateway Protocol (BGP) configuration for neighbor groups, use the **show bgp neighbor-group** command in XR EXEC mode.

```
show bgp neighbor-group group-name {configuration [defaults] [nvgen] | inheritance | users}
```

Syntax Description

| | |
|----------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>group-name</i> | Name of the address family group to display. |
| configuration | (Optional) Displays the effective configuration for the neighbor group, including any configuration inherited by this neighbor group. |
| defaults | (Optional) Displays all configuration, including default configuration. |
| nvgen | (Optional) Displays output in show running-config command output. If the defaults keyword is also specified, the output is not suitable for cutting and pasting into a configuration session. |
| inheritance | Displays the af-groups, session groups, and neighbor groups from which this neighbor group inherits configuration. |
| users | Displays the neighbors and neighbor groups that inherit configuration from this neighbor group. |

Command Default

No default behavior or value

Command Modes

XR EXEC

Command History

| Release | Modification |
|---------------|------------------------------|
| 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.

Use the **show bgp neighbor-group** command with the *group-name* **configuration** argument and keyword to display the effective configuration of a neighbor group, including any configuration inherited from session groups, address family groups, and neighbor groups through application of the **use** command. The source of each configured command is also displayed.

Use the **defaults** keyword to display all configuration for the neighbor group, including default configuration. The command output identifies default onfiguration. Use the **nvgen** keyword to display configuration in the output form of **show running-config** command. Output in this form is suitable for cutting and pasting into a configuration session.

The **show bgp neighbor-group** command with the *group-name* **inheritance** argument and keyword displays the session groups, address family groups, and neighbor groups from which the specified neighbor group inherits configuration.

The **show bgp neighbor-group** *group-name* command displays the neighbors and neighbor groups that inherit configuration from the specified neighbor group.

| Task ID | Task ID | Operations |
|---------|---------|------------|
| | bgp | read |

Examples

The examples use the following configuration:

```
af-group group3 address-family ipv4 unicast
  remove-private-AS
  soft-reconfiguration inbound
!
af-group group2 address-family ipv4 unicast
  use af-group group3
  send-community-ebgp
  send-extended-community-ebgp
  capability orf prefix both
!
session-group group3
  dmzlink-bw
!
neighbor-group group3
  use session-group group3
  timers 30 90
!
neighbor-group group1
  remote-as 1982
  use neighbor-group group2
  address-family ipv4 unicast
!
!
neighbor-group group2
  use neighbor-group group3
  address-family ipv4 unicast
  use af-group group2
  weight 100
!
```

The following is sample output from the **show bgp neighbor-group** command with the **configuration** keyword:

```
RP/0/RP0/CPU0:router# show bgp neighbor-group group1 configuration

neighbor-group group1
  remote-as 1982                []
  timers 30 90                 [n:group2 n:group3]
  dmzlink-bw                   [n:group2 n:group3 s:group3]
  address-family ipv4 unicast  []
  capability orf prefix both  [n:group2 a:group2]
  remove-private-AS           [n:group2 a:group2 a:group3]
  send-community-ebgp         [n:group2 a:group2]
  send-extended-community-ebgp [n:group2 a:group2]
  soft-reconfiguration inbound [n:group2 a:group2 a:group3]
  weight 100                  [n:group2]
```

The configuration source is shown to the right of each command. In the output, the **remote-as** command is configured directly on neighbor group group1, and the **send-community-ebgp** command is inherited from neighbor group group2, which in turn inherits the setting from af-group group2.

The following is sample output from the **show bgp neighbor-group** command with the **users** keyword. This output shows that the group1 neighbor group inherits session (address family-independent configuration parameters) from the group2 neighbor group. The group1 neighbor group also inherits IPv4 unicast configuration parameters from the group2 neighbor group:

```
RP/0/RP0/CPU0:router# show bgp neighbor-group group2 users

Session:      n:group1
IPv4 Unicast: n:group1
```

The following is sample output from the **show bgp neighbor-group** command with the **inheritance** keyword. This output shows that the specified neighbor group group1 inherits session (address family-independent configuration) from neighbor group group2, which inherits its own session from neighbor group group3. Neighbor group group3 inherited its session from session group group3. It also shows that the group1 neighbor-group inherits IPv4 unicast configuration parameters from the group2 neighbor group, which in turn inherits them from the group2 af-group, which itself inherits them from the group3 af-group:

```
RP/0/RP0/CPU0:router# show bgp neighbor-group group1 inheritance

Session:      n:group2 n:group3 s:group3
IPv4 Unicast: n:group2 a:group2 a:group3
```

This table describes the significant fields shown in the display.

Table 15: 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 10 | Configures a BGP address family group. |
| session-group, on page 189 | Creates a session group and enters session group configuration mode. |

| Command | Description |
|-----------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------|
| show bgp af-group, on page 220 | Displays information about configuration for address family groups. |
| show bgp neighbors, on page 256 | Displays information about BGP neighbors, including configuration inherited from neighbor groups, session groups, and address family groups. |
| show bgp session-group, on page 316 | 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 368 | 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 XR 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]]
```

Syntax Description

| | |
|----------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| ipv4 | (Optional) Specifies IP Version 4 address prefixes. |
| unicast | (Optional) Specifies unicast 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. |
| 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. |
| <i>ip-address</i> | (Optional) IP address of the BGP-speaking neighbor. If you omit this argument, all neighbors are displayed. |
| advertised-routes | (Optional) Displays all routes the router advertised to the neighbor. |
| dampened-routes | (Optional) Displays the dampened routes that are learned from the neighbor. |
| flap-statistics | (Optional) Displays flap statistics of the routes learned from the neighbor. |
| received { prefix-filter routes } | (Optional) Displays information received from the BGP neighbor. The options are: prefix-filter — Displays the prefix list filter. routes —Displays routes from the neighbor before inbound policy |
| routes | (Optional) Displays routes learned from the neighbor. |
| configuration | (Optional) Displays the effective configuration for the neighbor, including any settings that have been inherited from session groups, neighbor groups, or af-groups used by this neighbor. |
| defaults | (Optional) Displays all configuration settings, including any default settings. |

| | |
|----------------------------|----------------------------------------------------------------------------------------------------------------------------------|
| nvgen | (Optional) Displays output in the show running-config command output. |
| inheritance | (Optional) Displays the session groups, neighbor groups, and af-groups from which this neighbor inherits configuration settings. |
| decoded-message-log | (Optional) Displays BGP message logs. |
| in | (Optional) Displays BGP inbound messages. |
| out | (Optional) Displays BGP outbound messages. |
| standby | (Optional) Displays information about the standby card. |

Command Default

If no address family or subaddress family is specified, the default address family and subaddress family specified using the **set default-afi** and **set default-safi** commands are used.

Command Modes

XR EXEC

Command History

| Release | Modification |
|---------------|------------------------------------------------------|
| Release 5.0.0 | This command was introduced. |
| Release 6.0 | 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 NCS 6000 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 12.12.12.5
...
  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 12.12.12.5 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/RP0/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

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%

```

This table describes the significant fields shown in the display.

Table 16: 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 | <ul style="list-style-type: none"> • 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) in asdot format is 1.0 to 65535.65535. |

| Field | Description |
|--------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| local AS | Autonomous system number of the local system. <ul style="list-style-type: none"> • 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. |
| internal link | Neighbor is an internal BGP peer. |
| external link | Neighbor is an external BGP peer. |
| Administratively shut down | Neighbor connection is disabled using the shutdown command. |
| remote router ID | Router ID (an IP address) of the neighbor. |
| Neighbor under common administration | Neighbor is internal or a confederation peer. |
| BGP state | Internal state of this BGP connection. |
| BFD enabled | Status of bidirectional forwarding detection. |
| TCP open mode | TCP mode used in establishing the BGP session. The following valid TCP mode are supported: <ul style="list-style-type: none"> • 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: <ul style="list-style-type: none"> • 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. |

| Field | Description |
|--------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Address family | Indicates that the local system supports the displayed address family capability. If “received” is displayed, the neighbor also supports the displayed address family. |
| Received | Number of messages received from this neighbor, the number of notification messages received and processed from this neighbor, and the number of messages that have been received, but not yet processed. |
| Sent | Number of messages sent to this neighbor, the number of notification messages generated to be sent to this neighbor, and the number of messages queued to be sent to this neighbor. |
| Minimum time between advertisement runs | Advertisement interval (in seconds) for this neighbor. |
| For Address Family | Information that follows is specific to the displayed address family. |
| BGP neighbor version | Last version of the BGP database that was sent to the neighbor for the specified address family. |
| Update group | Update group to which the neighbor belongs. |
| Route reflector client | Indicates that the local system is acting as a route reflector for this neighbor. |
| 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. |

| Field | Description |
|-----------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| AF-dependant capabilities | BGP capabilities that are specific to a particular address family. The following valid AF-dependent BGP capabilities are supported: <ul style="list-style-type: none"> • route refresh capability • route refresh capability OLD value |
| Outbound Route Filter | Neighbor has the Outbound Route Filter (ORF) capability for the specified address family. Details of the capabilities supported are also shown: Send-mode—“advertised” is shown if the local system can send an outbound route filter to the neighbor. “received” is shown if the neighbor can send an outbound route filter to the local system. Receive-mode—“advertised” is shown if the local system can receive an outbound route filter from the neighbor. “received” is shown if the neighbor can receive an outbound route filter from the local system. |
| Graceful Restart Capability | Indicates whether graceful restart capability has been advertised to and received from the neighbor for the specified address family. |
| Neighbor preserved the forwarding state during latest restart | Indicates that when the neighbor connection was last established, the neighbor indicated that it preserved its forwarding state for the specified address family. |
| 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. |

| Field | Description |
|-----------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Outbound path policy | Indicates if an outbound path policy is configured. |
| Incoming update prefix filter list | Indicates a prefix list is configured to filter inbound updates from the neighbor. |
| Default weight | Default weight for routes received from the neighbor. |
| Policy for incoming advertisements | Indicates a route policy is configured to be applied to inbound updates from the neighbor. |
| Policy for outgoing advertisements | Indicates a route policy is configured to be applied to outbound updates to the neighbor. |
| Type | Indicates whether the condition map selects routes that should be advertised, or routes that should not be advertised: Exist—Routes advertised if permitted by the condition route map. Non-exist—Routes advertised if denied by the condition route map. |
| accepted prefixes | Number of prefixes accepted. |
| Prefix advertised | Number of prefixes advertised to the neighbor during the lifetime of the current connection with the neighbor. |
| 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. |

| Field | Description |
|-----------------------------------------------------|---------------------------------------------------------------------------------------------------|
| Time since last notification received from neighbor | Amount of time since a notification message was last received from the neighbor. |
| Error Code | Type of notification that was received. The notification data received, if any, is also displayed |
| External BGP neighbor may be up to <n> hops away | Indicates ebgp-multihop is configured for the neighbor. |
| External BGP neighbor not directly connected | Indicates that the neighbor is not directly attached to the local system. |
| Notification data sent: | Data providing more details on the error along with the error notification sent to the neighbor. |

The following is sample output from the **show bgp neighbors** command with the **advertised-routes** keyword:

```
RP/0/RP0/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/RP0/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 17: show bgp neighbors routes Field Descriptions

| Field | Description |
|--------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| BGP router identifier | BGP identifier for the local system. |
| local AS number | Autonomous system number for the local system. |
| BGP main routing table version | Last version of the BGP database that was installed into the main routing table. |
| Dampening enabled | Displayed if dampening is enabled for the routes in this BGP routing table. |
| BGP scan interval | Interval (in seconds) between scans of the BGP table specified by the address family and subaddress family. |
| Status codes | <p>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):</p> <p>S—Path is stale, indicating that a graceful restart is in progress with the peer from which the route was learned.</p> <p>s—Path is more specific than a locally sourced aggregate route and has been suppressed.</p> <p>*—Path is valid.</p> <p>The second character may be (in order of precedence):</p> <p>>—Path is the best path to use for that network.</p> <p>d—Path is dampened.</p> <p>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.</p> <p>The third character may be:</p> <p>i—Path was learned by an internal BGP (iBGP) session.</p> |
| Origin codes | <p>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:</p> <p>i—Path originated from an Interior Gateway Protocol (IGP) and was advertised with a network or aggregate-address command.</p> <p>e—Path originated from an Exterior Gateway Protocol (EGP).</p> <p>?—Origin of the path is not clear. Usually, this is a route that is redistributed into BGP from an IGP.</p> |
| 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. |

The following is sample output from the **show bgp neighbors** command with the **dampened-routes** keyword:

```
RP/0/RP0/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 18: 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. |

| Field | Description |
|-------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| BGP scan interval | Interval (in seconds) between scans of the BGP table specified by the address family and subaddress family. |
| Status codes | <p>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):</p> <p>S—Path is stale, indicating that a graceful restart is in progress with the peer from which the route was learned.</p> <p>s—Path is more specific than a locally sourced aggregate route and has been suppressed.</p> <p>*—Path is valid.</p> <p>The second character may be (in order of precedence):</p> <p>>—Path is the best path to use for that network.</p> <p>d—Path is dampened.</p> <p>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.</p> <p>The third character may be:</p> <p>i—Path was learned by an internal BGP (iBGP) session.</p> |
| Origin codes | <p>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:</p> <p>i—Path originated from an Interior Gateway Protocol (IGP) and was advertised with a network or aggregate-address command.</p> <p>e—Path originated from an Exterior Gateway Protocol (EGP).</p> <p>?—Origin of the path is not clear. Usually, this is a route that is redistributed into BGP from an IGP.</p> |
| 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/RP0/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 19: 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 | <p>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):</p> <ul style="list-style-type: none"> 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. <p>The second character may be (in order of precedence):</p> <ul style="list-style-type: none"> 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. <p>The third character may be:</p> <ul style="list-style-type: none"> 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 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. |
| 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/RP0/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 20: 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. |

| Field | Description |
|------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 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. |
| 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/RP0/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 []

  policy pass-all in []
  policy pass-all out []
```

This table describes the significant fields shown in the display.

Table 21: 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/RP0/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/RP0/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 []
```

Related Commands

| Command | Description |
|-------------------------------------------------------|---------------------------------------------------------------------------------------------------------|
| clear bgp, on page 82 | Resets a BGP connection or session. |
| network (BGP), on page 143 | Specifies a local network that the BGP routing process should originate and advertise to its neighbors. |
| route-policy (BGP), on page 178 | 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 193 | Displays entries in the BGP routing table. |
| show bgp dampened-paths, on page 238 | Displays BGP dampened routes. |
| show bgp flap-statistics, on page 242 | Displays BGP routes that have flapped. |
| show bgp neighbor-group, on page 252 | Displays information about the BGP configuration for neighbor groups. |
| shutdown (BGP), on page 339 | 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 XR EXEC mode.

```
show bgp [{ipv4 {unicast | all} | ipv6 {unicast | all} | unicast | unicast}] 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. |
| all | (Optional) For address family, specifies prefixes for all address families. |
| standby | (Optional) Displays information about the standby card. |

Command Default No default behavior or values.

Command Modes XR EXEC

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | Release 5.2.1 | 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 |
|---------|---------|------------|
| | bgp | read |

Examples

The following is sample output from the **show bgp neighbors nsr** command with the **standby** keyword:

```
RP/0/RP0/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 22: 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 155 | Activates the Border Gateway Protocol (BGP) nonstop routing (NSR). |
| show bgp summary nsr, on page 324 | Displays the Border Gateway Protocol (BGP) nonstop routing (NSR) information. |
| show bgp summary, on page 320 | 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 XR EXEC mode.

```
show bgp nexthops A.B.C.D.aigp-value[statistics] [speaker speaker-id] [standby]
```

| Syntax Description | | |
|----------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------|
| ipv4 | (Optional) | Specifies IP Version 4 address prefixes. |
| unicast | (Optional) | Specifies unicast address prefixes. |
| 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. |
| statistics | (Optional) | Specifies nexthop statistics. |
| speaker <i>speaker-id</i> | (Optional) | Specifies a speaker process ID. |
| standby | (Optional) | Displays information about the standby card. |
| Command Default | No default behavior or value | |
| Command Modes | XR EXEC | |
| Command History | Release | Modification |
| | Release 5.0.0 | This command was introduced. |
| Usage Guidelines | <p>To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.</p> <p>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).</p> <p>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 <i>speaker-id</i> keyword and argument to display information for only the specified speaker process.</p> | |
| Task ID | Task ID | Operations |
| | bgp | read |

Related Commands

| Command | Description |
|-------------------------------------------------------|---------------------------------------------------------------|
| bgp redistribute-internal, on page 70 | 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 XR EXEC mode.

```
show bgp [{ipv4 {unicast | labeled-unicast | all} | ipv6 {unicast | all | labeled-unicast} | all {unicast | all | labeled-unicast} | unicast}] nsr [standby]
```

| Syntax Description | |
|------------------------|-------------------------------------------------------------------------------------|
| ipv4 | (Optional) Specifies IP Version 4 address prefixes. |
| unicast | (Optional) Specifies unicast 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. |
| 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 XR EXEC

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | Release 5.2.1 | 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 |
|---------|---------|------------|
| | bgp | read |

Examples

The following is sample output from the **show bgp nsr** command:

```
RP/0/RP0/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: IGP 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

```

```
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/RP0/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: IGP 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

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 155 | 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 XR EXEC mode.

```
show bgp paths [detail] [debug] [regexp regular-expression]
```

| Syntax Description | detail | (Optional) Displays detailed attribute information. |
|--------------------|-----------------------------------------|-------------------------------------------------------------------------------------------------|
| | debug | (Optional) Displays attribute process ID, hash bucket, and hash chain ID attribute information. |
| | regexp <i>regular-expression</i> | (Optional) Specifies an autonomous system path that matches the regular expression. |

Command Default No default behavior or values

Command Modes XR EXEC

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | 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.

Use the **show bgp paths** command to display information about AS paths and the associated attributes with which the paths were received.

If no options are specified, all stored AS paths are displayed with the number of routes using each path.



Note The AS path information is stored independently of the address family, making it possible that routes from different address families could be using the same path.

Use the *regular-expression* argument to limit the output to only those paths that match the specified regular expression. See the for information on regular expressions.

Use the **detail** keyword to display detailed information on the attributes stored with the AS path.

| Task ID | Task ID | Operations |
|---------|---------|------------|
| | bgp | read |

Examples

The following is sample output from the **show bgp paths** command:

```
RP/0/RP0/CPU0:router# show bgp paths detail
```

```

Proc  Attributes                               Refcount   Metric Path
Spk 0  ORG AS LOCAL                             7           0 i
Spk 0  ORG AS LOCAL COMM EXTCOMM             3           0 21 i
Spk 0  MET ORG AS                             3           55 2 i
Spk 0  ORG AS                                 3           0 2 10 11 i
Spk 0  ORG AS COMM                           3           0 2 10 11 i
Spk 0  MET ORG AS ATOM                       3           2 2 3 4 ?
Spk 0  MET ORG AS                             3           1 2 3 4 e
Spk 0  MET ORG AS                             3           0 2 3 4 i

```

This table describes the significant fields shown in the display.

Table 23: 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. <ul style="list-style-type: none"> • 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. |
| Refcount | Number of routes using a path. |
| Metric | Value of the interautonomous system metric, otherwise known as the MED metric. |
| Path | Autonomous system path to the destination network. At the end of the path is the origin code for the path: i—Path originated from an Interior Gateway Protocol (IGP) and was advertised with a network or aggregate-address command. e—Path originated from an Exterior Gateway Protocol (EGP). ?—Origin of the path is not clear. Usually, this is a route that is redistributed into BGP from an IGP. |

show bgp policy

To display information about Border Gateway Protocol (BGP) advertisements under a proposed policy, use the **show bgp policy** command in XR EXEC mode.

show bgp policy

| Syntax | Description |
|----------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------|
| ipv4 | (Optional) Specifies IP Version 4 address prefixes. |
| unicast | (Optional) Specifies unicast address prefixes. |
| 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. |
| <i>ip-address</i> | (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. |
| standby | (Optional) Displays information about the standby card. |
| summary | (Optional) Displays a summary of the BGP advertisements. |

Command Default Advertisements for all neighbors are displayed if the **neighbor** *ip-address* keyword and argument are not specified. If no address family or subaddress family is specified, the default address family and subaddress family specified using the **set default-afi** and **set default-safi** commands are used.

Command Modes XR EXEC

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | 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.



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 NCS 6000 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.

Task ID

| Task ID | Operations |
|---------|------------|
| bgp | read |

Examples

The following is sample output from the **show bgp policy** command with the **summary** keyword in XR EXEC mode:

```
RP/0/RP0/CPU0:router# show bgp policy summary

Network          Next Hop          From              Advertised to
172.16.1.0/24    10.0.101.1       10.0.101.1       10.0.101.2
                                                         10.0.101.3

172.17.0.0/16    0.0.0.0          Local             10.0.101.1
                                                         10.0.101.2
                                                         10.0.101.3
```

This table describes the significant fields shown in the display.

Table 24: show bgp policy summary Field Descriptions

| Field | Description |
|---------|--------------------------------------------|
| 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. |
| From | IP address of the peer that advertised this route. |
| Local | Indicates the route originated on the local system. |
| Local Aggregate | Indicates the route is an aggregate created on the local system. |
| Advertised to | Indicates the neighbors to which this route was advertised. |

The following is sample output from the **show bgp policy** command in XR EXEC mode:

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

11.0.0.0/24 is advertised to 10.4.101.1
  Path info:
    neighbor: Local           neighbor router id: 10.4.0.1
    valid local best
  Attributes after inbound policy was applied:
    next hop: 0.0.0.0
    MET ORG AS
    origin: IGP metric: 0
    aspath:
  Attributes after outbound policy was applied:
    next hop: 10.4.0.1
    MET ORG AS
    origin: IGP metric: 0
    aspath: 1

11.0.0.0/24 is advertised to 10.4.101.2
  Path info:
    neighbor: Local           neighbor router id: 10.4.0.1
    valid local best
  Attributes after inbound policy was applied:
    next hop: 0.0.0.0
    MET ORG AS
    origin: IGP metric: 0
    aspath:
  Attributes after outbound policy was applied:
    next hop: 10.4.0.1
    MET ORG AS
    origin: IGP metric: 0
    aspath:

11.0.0.0/24 is advertised to 10.4.101.3
  Path info:
    neighbor: Local           neighbor router id: 10.4.0.1
    valid local best
  Attributes after inbound policy was applied:
    next hop: 0.0.0.0
    MET ORG AS
    origin: IGP metric: 0
    aspath:
  Attributes after outbound policy was applied:
    next hop: 10.4.0.1
    MET ORG AS
    origin: IGP metric: 0
```

```

    aspath:

12.0.0.0/24 is advertised to 10.4.101.2
Path info:
  neighbor: 10.4.101.1      neighbor router id: 10.4.101.1
  valid external best
Attributes after inbound policy was applied:
  next hop: 10.4.101.1
  ORG AS
  origin: IGP neighbor as: 2
  aspath: 2 3 4
Attributes after outbound policy was applied:
  next hop: 10.4.101.1
  ORG AS
  origin: IGP neighbor as: 2
  aspath:2 3 4

12.0.0.0/24 is advertised to 10.4.101.3
Path info:
  neighbor: 10.4.101.1      neighbor router id: 10.4.101.1
  valid external best
Attributes after inbound policy was applied:
  next hop: 10.4.101.1
  ORG AS
  origin: IGP neighbor as: 2
  aspath: 2 3 4
Attributes after outbound policy was applied:
  next hop: 10.4.101.1
  ORG AS
  origin: IGP neighbor as: 2
  aspath:2 3 4

```

This table describes the significant fields shown in the display.

Table 25: 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. |
| 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. |

| Field | Description |
|-----------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| (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. |
| 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. |

| Field | Description |
|---------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Attributes after inbound policy was applied | <p>Displays attributes associated with the received route, after any inbound policy has been applied.</p> <p>AGG—Aggregator attribute is present.</p> <p>AS—AS path attribute is present.</p> <p>ATOM—Atomic aggregate attribute is present.</p> <p>COMM—Communities attribute is present.</p> <p>EXTCOMM—Extended communities attribute is present.</p> <p>LOCAL—Local preference attribute is present.</p> <p>MET—Multi Exit Discriminator (MED) attribute is present.</p> <p>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.</p> <p>ORG—Origin attribute is present.</p> |
| origin | <p>Origin of the path:</p> <p>IGP—Path originated from an Interior Gateway Protocol (IGP) and was sourced by BGP using a network or aggregate-address command.</p> <p>EGP—Path originated from an Exterior Gateway Protocol.</p> <p>incomplete—Origin of the path is not clear; in example, a route that is redistributed into BGP from an IGP.</p> |
| 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. |
| community | <p>Community attributes associated with the path. Community values are displayed in AA:NN format, except for the following well-known communities:</p> <p>Local-AS—Community with value 4294967043 or hex 0xFFFFFFFF03. Routes with this community value are not advertised outside the local autonomous system or confederation boundary.</p> <p>no-advertise—Community with value 4294967042 or hex 0xFFFFFFFF02. Routes with this community value are not advertised to any BGP peers.</p> <p>no-export—Community with value 4294967041 or hex 0xFFFFFFFF01. Routes with this community are not advertised to external BGP peers, even if those peers are in the same confederation as the local router.</p> |

| Field | Description |
|----------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 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 178 | Applies an inbound or outbound routing policy to a 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 advertised, on page 214 | Displays routes advertised to neighbors. |
| show bgp neighbors, on page 256 | Displays information about the TCP and BGP connections to neighbors. |
| show bgp route-policy, on page 312 | 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 XR EXEC mode.

```
show bgp process [performance-statistics] [detail] [standby]
```

| Syntax Description | |
|--------------------------------|------------------------------------------------------------------------------------------------|
| ipv4 | (Optional) Specifies IP Version 4. |
| unicast | (Optional) Specifies the unicast subaddress family. |
| 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. |
| all | (Optional) For address family, specifies prefixes for all address families. |
| performance- statistics | (Optional) Displays performance statistics relative to the work done by the specified process. |
| detail | (Optional) Specifies detailed process information. |
| standby | (Optional) Displays information about the standby card. |

Command Default If no address family or subaddress family is specified, the default address family and subaddress family specified using the **set default-afi** and **set default-safi** commands are used.

Command Modes XR EXEC

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | 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.



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 NCS 6000 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/RP0/CPU0:router# show bgp process

BGP Process Information
BGP is operating in STANDALONE mode
Autonomous System: 1
Router ID: 10.0.0.5 (manually configured)
Cluster ID: 10.0.0.5
Fast external fallover enabled
Neighbor logging is enabled
Enforce first AS enabled
Default local preference: 100
Default keepalive: 60
Update delay: 120
Generic scan interval: 60

Address family: IPv4 Unicast
Dampening is enabled
Client reflection is enabled
Scan interval: 60
Main Table Version: 150
IGP notification: IGP 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 26: show bgp process Field Descriptions

| Field | Description |
|---------------------|---------------------------------------------------------------------------------|
| BGP is operating in | Indicates BGP is operating in standalone mode. This is the only supported mode. |

| Field | Description |
|----------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Autonomous System | Autonomous system number for the local system. <ul style="list-style-type: none"> • 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. |
| Router ID | BGP identifier assigned to the local system. If this is explicitly configured using the bgp router-id command, “manually configured” is displayed. If the router ID is not manually configured, it is determined from a global router ID. If no global ID is available, the router ID is shown as 0.0.0.0. |
| Confederation ID | Confederation identifier for the local system. |
| Cluster ID | Cluster identifier for the local system. If this is manually configured using the bgp cluster-id command, “manually configured” is displayed. |
| Default metric | Default metric. This is controlled by the default-metric command. |
| Fast external fallover enabled | Indicates whether fast external fallover is enabled. This is controlled by the bgp fast-external-fallover disable command. |
| Neighbor logging enabled | Indicates whether logging of peer connection up and down transitions is enabled. This is controlled by the bgp log neighbor changes disable command. |
| Enforce first AS enabled | Indicates that strict checking of the first AS number in paths received from external BGP peers is enabled. This is controlled by the bgp enforce-first-as disable command. |
| iBGP to IGP redistribution | Indicates internal redistribution is enabled using the bgp redistribution-internal command. |
| Treating missing MED as worst | Indicates missing Multi Exit Discriminator (MED) metric values are treated as worst in the route selection algorithm. This is controlled by the bgp bestpath med missing-as-worst command. |
| Always compare MED is enabled | Indicates that the MED is always used during the route selection algorithm, even when paths are received from external BGP neighbors in different autonomous systems. This is controlled by the bgp bestpath med always command. |
| AS Path ignore is enabled | Indicates that the AS path length is ignored by the route selection algorithm. This is controlled by the bgp bestpath as-path ignore command. |
| Comparing MED from confederation peers | Indicates that the MED values are used in the route selection algorithm when comparing routes received from confederation peers. This is controlled by the bgp bestpath med confed command. |

| Field | Description |
|------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Comparing router ID for eBGP paths | Indicates that the router ID is used as a tiebreaker by the route selection algorithm when comparing identical routes received from different external BGP neighbors. This is controlled by the bgp bestpath compare-routerid command. |
| Default local preference | Default local preference value used for BGP routes. 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: <ul style="list-style-type: none"> • bgp graceful-restart • bgp graceful-restart purge-time • bgp graceful-restart stalepath-time • bgp graceful-restart restart-time • bgp graceful-restart graceful-reset |
| Update delay | Maximum time that a BGP process stays in read-only mode. |
| Generic scan interval | Interval (in seconds) between BGP scans for address family-independent tasks. This is controlled by the bgp scan-time command. |
| Dampening | Indicates whether dampening is enabled for the specified address family. This is controlled by the dampening command. |
| Client reflection | Indicates whether client-to-client route reflection is enabled for the specified address family. This is controlled by the bgp client-to-client reflection disable command. |
| Scan interval | Interval (in seconds) between BGP scans for the given address family. This is controlled by the bgp scan-time command in address family configuration mode. |
| Main Table Version | Last version of the BGP database that was installed into the main routing table. |
| IGP notification | Indicates whether Interior Gateway Protocols (IGP) have been notified of BGP convergence for the specified address family. |
| Node | Node on which the process is executing. |
| Process | Type of BGP process. |
| Speaker | Speaker process. A speaker process is responsible for receiving, processing, and sending BGP messages to configured neighbors. |
| Nbrs | Number of neighbors for which the process is responsible. |
| Estab | Number of neighbors that have connections in the established state for this process. |

| Field | Description |
|----------|----------------------------------------------------------|
| Rst | Number of times this process was restarted. |
| Upd-Rcvd | Number of update messages received by the process. |
| Upd-Sent | Number of update messages sent by the process. |
| Nfn-Rcvd | Number of notification messages received by the process. |
| Nfn-Sent | Number of notification messages sent by the process. |

The following is sample output from the **show bgp process** command with the Graceful Maintenance feature enabled:

```
RP/0/0/CPU0:R1#show bgp process
```

```
...
Graceful Maintenance active. Retaining routes in RIB during BGP shutdown
...
```

Or

```
Graceful Maintenance active for all neighbors. Retaining routes in RIB during BGP shutdown
```

```
*****
```

```
RP/0/0/CPU0:Jan 28 22:01:36.356 : bgp[1056]: %ROUTING-BGP-5-ADJCHANGE : neighbor 10.10.10.4
Up (VRF: default) (AS: 4) WARNING: Graceful Maintenance is Active
```

The following is sample output from the **show bgp process** command with the **detail** keyword:

```
RP/0/RP0/CPU0:router# show bgp all all process detail
```

```
BGP Process Information
BGP is operating in STANDALONE mode
Autonomous System: 1
Router ID: 10.0.0.5 (manually configured)
Cluster ID: 10.0.0.5
Fast external fallover enabled
Neighbor logging is enabled
Enforce first AS enabled
Default local preference: 100
Default keepalive: 60
Update delay: 120
Generic scan interval: 60
```

```
BGP Speaker process: 0, location node0_0_0
Neighbors: 3, established: 2
```

```

                Sent           Received
Updates:         3             15
Notifications:  0             0
```

```

                Number         Memory Used
Attributes:     12            1104
AS Paths:       10            400
Communities:    2            1080
Extended communities: 1         40
Route Reflector Entries: 0         0
Route-map Cache Entries: 0         0
Filter-list Cache Entries: 0        0
```

```

Next Hop Cache Entries:    2                80
Update messages queued:   0

Address family: IPv4 Unicast
Dampening is enabled
Client reflection is enabled
Main Table Version: 12
IGP notification: IGP notified

State: normal mode.
BGP Table Version: 12
Network Entries: 15, Soft Reconfig Entries: 0
Dampened Paths: 0, History Paths: 9

Prefixes:                  Allocated      Freed
Paths:                     15           0
                           19           0

Prefixes:                  Number       Memory Used
Paths:                     15          1230
                           19          760

```

This table describes the significant fields shown in the display.

Table 27: show bgp process detail Field Descriptions

| Field | Description |
|--------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| BGP is operating in | Indicates whether BGP is operating in standalone mode. |
| Autonomous System | Autonomous system number for the local system. |
| Router ID | BGP identifier assigned to the local system. If this is explicitly configured using the bgp router-id command, “manually configured” is displayed. If the router ID is not manually configured, it is determined from a global router ID. If the global ID is not available, the router ID is shown as 0.0.0.0. |
| Confederation ID | Confederation identifier for the local system. |
| Cluster ID | Cluster identifier for the local system. If this is manually configured using the bgp cluster-id command, “manually configured” is displayed. |
| Default metric | Default metric. |
| Fast external fallover enabled | Indicates whether fast external fallover is enabled. |
| Neighbor logging enabled | Indicates whether logging of peer connection up and down transitions is enabled. |
| Enforce first AS enabled | Indicates that strict checking of the first autonomous system (AS) number in paths received from external BGP peers is enabled. |
| 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. |

| Field | Description |
|----------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 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: <ul style="list-style-type: none"> • bgp graceful-restart • bgp graceful-restart purge-time • bgp graceful-restart stalepath-time • bgp graceful-restart restart-time • bgp graceful-restart graceful-reset |
| Update delay | Maximum time that a BGP process stays in read-only mode. |
| Generic scan interval | Interval (in seconds) between BGP scans for address family-independent tasks. This is controlled by the bgp scan-time command. |
| BGP Speaker Process | Speaker process responsible for receiving, processing and sending BGP messages. |
| Node | Node on which the specified process is executing. |
| Neighbors | Number of neighbors for which the specified process is responsible. |
| established | Number of neighbors that have connections in the established state for the specified process. |
| 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. |

| Field | Description |
|-------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 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 IGP's 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 | <p>BGP system state for the specified address family and process. This may be one of the following:</p> <p>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.</p> <p>best-path calculation mode—Route selection is being performed for the routes that were received while in read-only mode.</p> <p>RIB update mode—Routes that were selected in best-path calculation mode are being installed in the global RIB.</p> <p>label allocation mode: Labels are allocated for the received prefixes based on the requirement.</p> <p>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.</p> |
| BGP Table Version | Last version used in the BGP database for received routes. |
| Attribute download | Indicates whether the RIB attribute download is enabled. |
| Network Entries | Number of sets of prefix information held in the specified BGP process for the specified address family. |
| Soft Reconfig Entries | Number of sets of prefix information that are present only for the purpose of supporting soft reconfiguration. |
| Dampened Paths | Number of routes that are suppressed due to dampening for the specified address family. |
| History Paths | Number of routes that are currently withdrawn, but are being maintained to preserve dampening information. |
| Prefixes (Allocated/Freed) | Number of sets of prefix information for the specified address family that have been allocated and freed during the lifetime of the process. |
| Paths (Allocated/Freed) | Number of sets of route information for the specified address family that have been allocated and freed during the lifetime of the process. |
| Prefixes (Number/Memory Used) | Number of sets of prefix information currently allocated for the specified address family, and the amount of memory used by them. |
| Paths (Number/Memory Used) | Number of sets of route information currently allocated for the specified address family, and the amount of memory used by them. |

The following is sample output from the **show bgp process** command with the **performance-statistics** keyword:

```
RP/0/RP0/CPU0:router# show bgp process performance-statistics detail
BGP Speaker process: 0, Node: node0_0_CPU0
```

Restart count: 2
 Neighbors: 3, established: 2

| | Sent | Received |
|----------------|------|----------|
| Updates: | 20 | 20 |
| Notifications: | 0 | 0 |

| | Number | Memory Used |
|----------------------------|--------|-------------|
| Attributes: | 2 | 184 |
| AS Paths: | 2 | 48 |
| Communities: | 0 | 0 |
| Extended communities: | 0 | 0 |
| Route Reflector Entries: | 0 | 0 |
| Route-map Cache Entries: | 0 | 0 |
| Filter-list Cache Entries: | 0 | 0 |
| Next Hop Cache Entries: | 2 | 80 |
| Update messages queued: | 0 | |

Read 14 messages (1142 bytes) in 12 calls (time spent: 0.024 secs)
 Read throttled 0 times
 Processed 14 inbound messages (time spent: 0.132 secs)
 Wrote 2186 bytes in 24 calls (time spent: 0.024 secs)
 Processing write list: wrote 18 messages in 4 calls (time spent: 0.000 secs)
 Processing write queue: wrote 10 messages in 20 calls (time spent: 0.000 secs)
 Socket setup (LPTS): 4 calls (time spent: 0.010 secs)
 Configuration: 1 requests (time spent: 0.002 secs)
 Operational data: 9 requests (time spent: 0.026 secs)

State: normal mode.
 BGP Table Version: 150
 Network Entries: 149, Soft Reconfig Entries: 0

| | Allocated | Freed |
|-----------|-----------|-------|
| Prefixes: | 149 | 0 |
| Paths: | 200 | 0 |

| | Number | Memory Used |
|-----------|--------|-------------|
| Prefixes: | 149 | 12516 |
| Paths: | 200 | 8000 |

Updates generated: 149 prefixes in 8 messages from 2 calls (time spent: 0.046 secs)
 Scanner: 2 scanner runs (time spent: 0.008 secs)
 RIB update: 1 rib update runs, 149 prefixes installed (time spent: 0.024 secs)
 Process has converged for IPv4 Unicast.

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 28: 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 IGP's 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. |
| Conv | Indicates whether the process has converged after the initial update. |

| Field | Description |
|-----------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------|
| 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 <i>n</i> seconds | Indicates that BGP has reached initial convergence for the IPv4 unicast 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/RP0/CPU0:router# show bgp process performance-statistics detail

BGP Speaker process: 0, Node: node0_0_CPU0
Restart count: 2
Neighbors: 3, established: 2

                Sent                Received
Updates:         20                  20
Notifications:   0                   0

                Number                Memory Used
Attributes:      2                    184
AS Paths:        2                    48
Communities:     0                    0
Extended communities: 0                0
Route Reflector Entries: 0            0
Route-map Cache Entries: 0            0
Filter-list Cache Entries: 0          0
Next Hop Cache Entries: 2             80
Update messages queued: 0

Read 14 messages (1142 bytes) in 12 calls (time spent: 0.024 secs)
Read throttled 0 times
Processed 14 inbound messages (time spent: 0.132 secs)
Wrote 2186 bytes in 24 calls (time spent: 0.024 secs)
Processing write list: wrote 18 messages in 4 calls (time spent: 0.000 secs)
Processing write queue: wrote 10 messages in 20 calls (time spent: 0.000 secs)
Socket setup (LPTS): 4 calls (time spent: 0.010 secs)
Configuration: 1 requests (time spent: 0.002 secs)
Operational data: 9 requests (time spent: 0.026 secs)

State: normal mode.
BGP Table Version: 150
Network Entries: 149, Soft Reconfig Entries: 0

                Allocated                Freed
```

```

Prefixes:                149            0
Paths:                   200            0

                          Number          Memory Used
Prefixes:                149            12516
Paths:                   200            8000

Updates generated: 149 prefixes in 8 messages from 2 calls (time spent: 0.046 secs)
Scanner: 2 scanner runs (time spent: 0.008 secs)
RIB update: 1 rib update runs, 149 prefixes installed (time spent: 0.024 secs)
Process has converged for IPv4 Unicast.

First neighbor established: 1082604050s
Entered DO_BESTPATH mode: 1082604055s
Entered DO_RIBUPD mode: 1082604055s
Entered Normal mode: 1082604055s
Latest UPDATE sent: 1082604056s

```

This table describes the significant fields shown in the display.

Table 29: 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. |
| 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. |

| Field | Description |
|------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 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 | <p>BGP system state for the specified address family and process. This may be one of the following:</p> <p>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.</p> <p>best-path calculation mode—Route selection is being performed for the routes that were received while in read-only mode.</p> <p>RIB update mode—Routes that were selected in best-path calculation mode are being installed in the global RIB.</p> <p>label allocation mode: Labels are allocated for the received prefixes based on the requirement.</p> <p>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.</p> |
| 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. |
| Scanner | Number of times the scanner has run for the specified address family and real time spent in scanner processing. |

| Field | Description |
|----------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 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 show bgp process detail command with information on additional paths send and receive information:

```

BGP Process Information:
BGP is operating in STANDALONE mode
Autonomous System number format: ASDOT
Autonomous System: 100
Router ID: 22.22.22.22 (manually configured)
Default Cluster ID: 2.2.2.2 (manually configured)
Active Cluster IDs: 2.2.2.2
Fast external fallover enabled
Neighbor logging is enabled
Enforce first AS enabled
AS Path multipath-relax is enabled
Default local preference: 100
Default keepalive: 60
Graceful restart enabled
Restart time: 120
Stale path timeout time: 360
RIB purge timeout time: 600
Non-stop routing is enabled
Update delay: 120
Generic scan interval: 60

.....
.....
Prefixes:           Allocated      Freed
Paths:             12              0
Path-elems:        60              0

                   Number          Mem Used
Prefixes:          12             1200
Paths:             60             3120
Path-elems:        12             624

```

| Related Commands | Command | Description |
|------------------|---------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------|
| | bgp bestpath as-path ignore, on page 30 | Sets the autonomous system path length to ignore when calculating preferred paths. |
| | bgp bestpath compare-routerid, on page 32 | 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 34 | Compare the Multi Exit Discriminator (MED) for paths from neighbors in different autonomous systems. |
| | bgp bestpath med missing-as-worst, on page 38 | Assume paths with no MED attribute have the most undesirable MED value possible when performing path selection. |
| | bgp cluster-id, on page 46 | Enables reflection of routes between route reflector clients using a BGP route reflector. |
| | bgp cluster-id, on page 46 | Configure the cluster ID if the BGP cluster has more than one route reflector. |
| | bgp default local-preference, on page 54 | Sets the default local preference value. |
| | bgp redistribute-internal, on page 70 | 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 72 | Configures a fixed router ID for a BGP-speaking router. |
| | default-metric (BGP), on page 108 | 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 73 | Configures scanning intervals. |
| | timers bgp, on page 353 | Sets default BGP timers. |

show bgp regexp

To display routes matching the autonomous system path regular expression, use the **show bgp regexp** command in XR EXEC mode.

show bgp regexp *regular-expression*

| Syntax Description | | |
|--------------------|---------------------------|-------------------------------------------------------------------------------------|
| | ipv4 | (Optional) Specifies IP Version 4 address prefixes. |
| | unicast | (Optional) Specifies unicast address prefixes. |
| | 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. |
| | <i>regular-expression</i> | Regular expression to match the BGP autonomous system paths. |

Command Default If no address family or subaddress family is specified, the default address family and subaddress family specified using the **set default-afi** and **set default-safi** commands are used.

Command Modes XR EXEC

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | 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.



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 NCS 6000 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:

```
RP/0/RP0/CPU0:router# show bgp regexp "^3 "

BGP router identifier 10.0.0.5, local AS number 1
BGP main routing table version 64
BGP scan interval 60 secs
Status codes: s suppressed, d damped, h history, * valid, > best
                i - internal, S stale
Origin codes: i - IGP, e - EGP, ? - incomplete
   Network        Next Hop           Metric LocPrf Weight Path
*>i172.20.17.121   10.0.101.2             100      0 3 2000 3000 i
*>i10.0.0.0        10.0.101.2             100      0 3 100 1000 i
*>i172.5.23.0/24   10.0.101.2             100      0 3 4 60 4378 i
```

This table describes the significant fields shown in the display.

Table 30: 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. |
| 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. |

| Field | Description |
|--------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Status codes | <p>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):</p> <p>S—Path is stale, indicating that a graceful restart is in progress with the peer from which the route was learned.</p> <p>s—Path is more specific than a locally sourced aggregate route and has been suppressed.</p> <p>*—Path is valid.</p> <p>The second character may be (in order of precedence):</p> <p>>—Path is the best path to use for that network.</p> <p>d—Path is dampened.</p> <p>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.</p> <p>The third character may be:</p> <p>i—Path was learned by an internal BGP (iBGP) session.</p> |
| Origin codes | <p>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:</p> <p>i—Path originated from an Interior Gateway Protocol (IGP) and was advertised with a network or aggregate-address command.</p> <p>e—Path originated from an Exterior Gateway Protocol (EGP).</p> <p>?—Origin of the path is not clear. Usually, this is a route that is redistributed into BGP from an IGP.</p> |
| 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. |
| 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 193 | Displays entries in the BGP routing table. |
| show bgp route-policy, on page 312 | 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 XR EXEC mode.

show bgp route-policy *route-policy-name* [**standby**]

| Syntax Description | | |
|--------------------------|------------|--------------------------------------------------------------------------|
| ipv4 | (Optional) | Specifies IP Version 4 address prefixes. |
| unicast | (Optional) | Specifies unicast address prefixes. |
| 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. |
| <i>route-policy-name</i> | | Name of a route policy. |
| standby | (Optional) | Displays information about the standby card. |

Command Default If no address family or subaddress family is specified, the default address family and subaddress family specified using the **set default-afi** and **set default-safi** commands are used.

Command Modes XR EXEC

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | 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.



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 NCS 6000 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 XR EXEC mode:

```
RP/0/RP0/CPU0:router# show bgp route-policy p1

BGP router identifier 172.20.1.1, local AS number 1820
BGP main routing table version 729
Dampening enabled
BGP scan interval 60 secs
Status codes: s suppressed, d damped, h history, * valid, > best
                i - internal, S stale
Origin codes: i - IGP, e - EGP, ? - incomplete
   Network          Next Hop           Metric LocPrf Weight Path
*  10.13.0.0/16      192.168.40.24             0 1878 704 701 200 ?
*  10.16.0.0/16      192.168.40.24             0 1878 704 701 i
```

This table describes the significant fields shown in the display.

Table 31: 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. |

| Field | Description |
|--------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Status codes | <p>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):</p> <p>S—Path is stale, indicating that a graceful restart is in progress with the peer from which the route was learned.</p> <p>s—Path is more specific than a locally sourced aggregate route and has been suppressed.</p> <p>*—Path is valid.</p> <p>The second character may be (in order of precedence):</p> <p>>—Path is the best path to use for that network.</p> <p>d—Path is dampened.</p> <p>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.</p> <p>The third character may be:</p> <p>i—Path was learned by an internal BGP (iBGP) session.</p> |
| Origin codes | <p>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:</p> <p>i—Path originated from an Interior Gateway Protocol (IGP) and was advertised with a network or aggregate-address command.</p> <p>e—Path originated from an Exterior Gateway Protocol (EGP).</p> <p>?—Origin of the path is not clear. Usually, this is a route that is redistributed into BGP from an IGP.</p> |
| 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 12 | Configures an aggregate entry in a BGP routing table. |
| network (BGP), on page 143 | Specifies a local network that the BGP routing process should originate and advertise to its neighbors. |
| route-policy (BGP), on page 178 | 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 283 | 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 XR EXEC mode.

```
show bgp session-group group-name {configuration [defaults] [nvgen] | inheritance | users}
```

Syntax Description

| | |
|----------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>group-name</i> | Name of the session family group to display. |
| configuration | (Optional) Displays the effective configuration for the session group, including any inherited configuration. |
| defaults | (Optional) Displays all configuration, including default configuration. |
| nvgen | (Optional) Displays output in the form of the show running-config command. If the defaults keyword also is specified, the output is not suitable for cutting and pasting into a configuration session. |
| inheritance | (Optional) Displays the session groups from which this session group inherits configuration. |
| users | (Optional) Display the session groups, neighbor groups, and neighbors that inherit configuration from this session group. |

Command Default

No default behavior or value

Command Modes

XR EXEC

Command History

| Release | Modification |
|---------------|------------------------------|
| 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.

Use the **show bgp session-group** command with the *group-name* **configuration** argument and keyword to display the effective configuration of a session group, including any configuration inherited from other session groups through application of the **use** command. The source for each configured command is also displayed.

Use the **defaults** keyword to display the value of all configuration, including default configuration. Use the **nvgen** keyword to display configuration in the form of the **show running-config** command output. Output in this form is suitable for cutting and pasting into a configuration session.

Use the **show bgp session-group** command with the *group-name* **inheritance** argument and keyword to display the session groups from which the specified session group inherits configuration.

Use the **show bgp session-group** command with the *group-name* **users** argument and keyword to display the neighbors, neighbor groups, and session groups that inherit configuration from the specified session group.

| Task ID | Task ID | Operations |
|---------|---------|------------|
| | bgp | read |

Examples

For the example shown here, the following configuration is used:

```
session-group group3
  advertisement-interval 5
  dmzlink-bw
  !
session-group group1
  use session-group group2
  update-source Loopback0
  !
session-group group2
  use session-group group3
  ebgp-multihop 2
```

The following example shows the **show bgp session-group** command with the **configuration** keyword:

```
RP/0/RP0/CPU0:router# show bgp session-group group1 configuration

session-group group1
  advertisement-interval 5[s:group2 s:group3]
  ebgp-multihop 2 [s:group2]
  update-source Loopback0 []
  dmzlink-bandwidth [s:group2 s:group3]
```

The source of each command is shown to the right of the command. For example, **update-source** is configured directly on session group group1. The **dmzlink-bandwidth** command is inherited from session group group2, which in turn inherits it from session group group3.

The following example shows the **show bgp session-group** command with the **users** keyword:

```
RP/0/RP0/CPU0:router# show bgp session-group group2 users

IPv4 Unicast:a:group1
```

The following example shows the **show bgp session-group** command with the **inheritance** keyword.

```
RP/0/RP0/CPU0:router# show bgp session-group group1 inheritance

Session:s:group2 s:group3
```

The command output shows that the session group group1 directly uses the group2 session group. The group2 session group uses the group3 session group.

This table describes the significant fields shown in the display.

Table 32: 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 189 | Configures a BGP session group. |
| show bgp neighbor-group, on page 252 | Displays information about the BGP configuration for neighbor groups. |
| show bgp neighbors, on page 256 | Displays information about BGP connections to neighbors. |

show bgp sessions

To display brief information about BGP neighbors, use the **show bgp sessions** command in XR EXEC mode.

```
show bgp sessions [not-established] [not-nsr-ready]
```

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

XR EXEC

Command History

| Release | Modification |
|---------------|------------------------------|
| 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.

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 |

Related Commands

| Command | Description |
|-------------------------------------------------|------------------------------------------------------------------------------------|
| show bgp neighbors, on page 256 | 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 XR EXEC mode.

```
show bgp [{ipv4 {unicast | labeled-unicast | all} | ipv6 {unicast | all | labeled-unicast}}] summary
[standby]
```

| Syntax Description | | |
|------------------------|------------|--------------------------------------------------------------------------|
| ipv4 | (Optional) | Specifies IP Version 4 address prefixes. |
| unicast | (Optional) | Specifies unicast 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. |
| standby | (Optional) | Displays information about the standby card. |

Command Default If no address family or subaddress family is specified, the default address family and subaddress family specified using the **set default-afi** and **set default-safi** commands are used.

Command Modes XR EXEC

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | 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.



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 NCS 6000 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 (RecvTblVer, bRIB/RIB, SendTblVer, and TblVer) are specific to the specified address family and subaddress family. All other information is global.

The table versions provide an indication of whether BGP is up to date with all work for the specified address family and subaddress family.

- bRIB/RIB < RecvTblVer—Some received routes have not yet been considered for installation in the global routing table.
- TblVer < SendTblVer—Some received routes have been installed in the global routing table but have not yet been considered for advertisement to this neighbor.

| Task ID | Task ID | Operations |
|---------|---------|------------|
| | bgp | read |

Examples

The following is sample output from the **show bgp summary** command:

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

BGP router identifier 10.0.0.0, local AS number 2
BGP generic scan interval 60 secs
BGP table state: Active
Table ID: 0xe0000000
BGP main routing table version 1
BGP scan interval 60 secs

BGP is operating in STANDALONE mode.

Process          RecvTblVer    bRIB/RIB    LabelVer    ImportVer    SendTblVer
Speaker          1             0           1           1           0

Neighbor        Spk   AS  MsgRcvd  MsgSent    TblVer    InQ  OutQ  Up/Down  St/PfxRcd
10.0.101.0      0     2     0        0         0        0    0    00:00:00 Idle
10.0.101.1      0     2     0        0         0        0    0    00:00:00 Idle
```

This table describes the significant fields shown in the display.

Table 33: 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 182 command. <ul style="list-style-type: none"> • 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. |

| Field | Description |
|-------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Table ID | BGP database identifier. |
| BGP main routing table version | Last version of the BGP database that was injected into the main routing table. |
| Dampening enabled | Displayed if dampening has been enabled for the routes in this BGP routing table. |
| BGP scan interval | Interval (in seconds) between scans of the BGP table specified by the address family and subaddress family. |
| BGP is operating in | Specifies BGP is operating in standalone mode. |
| Process | BGP process. |
| RecvTblVer | Last version used in the BGP database for received routes. |
| bRIB/RIB | Last version of the local BGP database that was injected into the main routing table. |
| LabelVer | Label version used in the BGP database for label allocation. |
| ImportVer | Last version of the local BGP database for importing routes. |
| SendTblVer | Latest version of the local BGP database that is ready to be advertised to neighbors. |
| Some configured eBGP neighbors do not have any policy | Some external neighbors exist that do not have both an inbound and outbound policy configured for every address family, using the route-policy (BGP) command. In this case, no prefixes are accepted and advertised to those neighbors. |
| Neighbor | IP address of a neighbor. |
| Spr | Speaker process that is responsible for the neighbor. Always 0. |
| AS | Autonomous system. |
| MsgRcvd | Number of BGP messages received from a neighbor. |
| MsgSent | Number of BGP messages sent to a neighbor. |
| 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. |

| Field | Description |
|-----------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| St/PfxRcd | <p>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.</p> <p>If the number of prefixes received exceeds the maximum allowed (as set by the maximum-prefix command), “(PfxRcd)” appears.</p> <p>If the connection has been shut down using the shutdown command, “(Admin)” appears.</p> <p>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.</p> <p>If the connection has been shut down due to out of memory (OOM), “(OOM)” appears.</p> |

Related Commands

| Command | Description |
|-------------------------------------------------|------------------------------------------------------------------------------------|
| route-policy (BGP), on page 178 | 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 XR EXEC mode.

show bgp summary nsr [standby]

| Syntax Description | |
|------------------------|-------------------------------------------------------------------------------------|
| ipv4 | (Optional) Specifies IP Version 4 address prefixes. |
| unicast | (Optional) Specifies unicast 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. |
| standby | (Optional) Displays information about the standby card. |

Command Default If no address family or subaddress family is specified, the default address family and subaddress family specified using the **set default-afi** and **set default-safi** commands are used.

Command Modes XR EXEC

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | Release 5.2.1 | 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 |
|---------|---------|------------|
| | bgp | read |

Examples

The following is sample output from the **show bgp summary nsr** command:

```
RP/0/RP0/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/RP0/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 34: 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. |
| 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 155 | Activates Border Gateway Protocol (BGP) nonstop routing (NSR) |
| show bgp nsr, on page 277 | 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 XR EXEC mode.

show bgp table

Syntax Description

ipv4 (Optional) Specifies IP Version 4 address prefixes.

unicast (Optional) Specifies unicast address prefixes.

ipv6 (Optional) Specifies IP Version 6 address prefixes.

Command Default

If no address family or subaddress family is specified, the default address family and subaddress family specified using the **set default-afi** and **set default-safi** commands are used.

Command Modes

XR EXEC

Command History

Release

Release 5.0.0

Modification

This command was introduced.

Usage Guidelines

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

Related Commands

| Command | Description |
|------------------------------------------------------|-------------------------------------------------------------------------------------------------|
| show bgp neighbor-group, on page 252 | Displays information about the Border Gateway Protocol (BGP) configuration for neighbor groups. |
| show bgp neighbors, on page 256 | Displays information about Border Gateway Protocol (BGP) connections to neighbors. |

| Command | Description |
|-----------------------------------------------|-----------------------------------------------------------------------|
| show bgp summary, on page 320 | 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 XR EXEC mode.

show bgptruncated-communities standby

| Syntax Description | Option | Description |
|--------------------|------------------------|-------------------------------------------------------------------------------------|
| | ipv4 | (Optional) Specifies IP Version 4 address prefixes. |
| | unicast | (Optional) Specifies unicast 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. |
| | standby | (Optional) Displays information about the standby card. |

Command Default If no address family or subaddress family is specified, the default address family and subaddress family specified using the **set default-afi** and **set default-safi** commands are used.

Command Modes XR EXEC

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | 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.



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 NCS 6000 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 aggregation or when inbound policy is applied.

| Task ID | Task ID | Operations |
|---------|---------|------------|
| | bgp | read |

Examples

The following is sample output from the **show bgp truncated-communities** command:

```
RP/0/RP0/CPU0:router# show bgp truncated-communities

BGP router identifier 172.20.1.1, local AS number 1820
BGP main routing table version 3042
BGP scan interval 60 secs
Status codes: s suppressed, d damped, h history, * valid, > best
               i - internal, S stale
Origin codes: i - IGP, e - EGP, ? - incomplete
   Network        Next Hop           Metric LocPrf Weight Path
*  10.13.0.0/16   192.168.40.24             0 1878 704 701 200 ?
*> 10.16.0.0/16   192.168.40.24             0 1878 704 701 i
```

This table describes the significant fields shown in the display.

Table 35: show bgp truncated-communities Field Descriptions

| Field | Description |
|--------------------------------|-------------------------------------------------------------------------------------------------------------|
| BGP router identifier | BGP Identifier for the local system. |
| local AS number | Autonomous system number for the local system. |
| BGP main routing table version | Last version of the BGP database that was installed into the main routing table. |
| 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 | <p>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):</p> <p>S—Path is stale, indicating that a graceful restart is in progress with the peer from which the route was learned.</p> <p>s—Path is more specific than a locally sourced aggregate route and has been suppressed.</p> <p>*—Path is valid.</p> <p>The second character may be (in order of precedence):</p> <p>>—Path is the best path to use for that network.</p> <p>d—Path is dampened.</p> <p>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.</p> <p>The third character may be:</p> <p>i—Path was learned by an internal BGP (iBGP) session.</p> |
| Origin codes | <p>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:</p> <p>i—Path originated from an Interior Gateway Protocol (IGP) and was advertised with a network or aggregate-address command.</p> <p>e—Path originated from an Exterior Gateway Protocol (EGP).</p> <p>?—Origin of the path is not clear. Usually, this is a route that is redistributed into BGP from an IGP.</p> |
| 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 12 | Creates an aggregate entry in a BGP routing table. |
| network (BGP), on page 143 | Specifies a local network that the BGP routing process should originate and advertise to its neighbors. |
| route-policy (BGP), on page 178 | 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 193 | 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 XR EXEC mode.

```
show bgp [{ipv4 {unicast | multicast | labeled-unicast | all | tunnel | mdt} | ipv6 {unicast | multicast | all | labeled-unicast} | all {unicast | multicast | all | labeled-unicast | mdt | tunnel} | vpnv4 unicast | vrf {vrf-name | all} [{ipv4 {unicast | labeled-unicast} | ipv6 unicast}] | vpnv6 unicast}] update-group [{neighbor ip-address | process-id . index [{summary | performance-statistics}]]]
show bgp [{ipv4 {unicast | labeled-unicast | all} | ipv6 {unicast | all | labeled-unicast}}] update-group [{neighbor ip-address | process-id . index [{summary | performance-statistics}]]] [standby]
```

| Syntax Description | | |
|-----------------------------------|------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| ipv4 | (Optional) | Specifies IP Version 4 update groups. |
| unicast | (Optional) | Specifies unicast update groups. |
| labeled-unicast | (Optional) | Specifies labeled unicast address prefixes. |
| all | (Optional) | Displays both unicast update groups. |
| ipv6 | (Optional) | Specifies IP Version 6 update groups. |
| all | (Optional) | Displays both IP Version 4 and IP Version 6 update groups. |
| neighbor <i>ip-address</i> | (Optional) | Specifies information on an update group for a specific neighbor. |
| <i>process-id.index</i> | (Optional) | Update group index. Process ID range is 0 to 254. Index range is 0 to 4294967295. The <i>process id.index</i> 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. |
| standby | (Optional) | Displays information about the standby card. |

Command Default If no address family or subaddress family is specified, the default address family and subaddress family specified using the **set default-afi** and **set default-safi** commands are used.

Command Modes XR EXECEXEC

| Command History | Release | Modification |
|-----------------|---------------|--------------------------------------------------------------------------------------------------------------|
| | Release 5.0.0 | This command was introduced. |
| | 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 NCS 6000 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 320](#) 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 |

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 134 | Limits the number of prefixes that can be received from a neighbor. |

| Command | Description |
|-------------------------------------------------|------------------------------------------------------------------------------------|
| route-policy (BGP), on page 178 | 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 320 | Displays the status of all BGP connections. |
| shutdown (BGP), on page 339 | Disables a neighbor without removing its configuration. |

show protocols (BGP)

To display information about the Border Gateway Protocol (BGP) instances running on the router, use the **show protocols** command in XR EXEC mode and specify either the **bgp** or **all** keyword.

```
show protocols [{ipv4 | ipv6 | afi-all}] [{allprotocol}]
```

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

XR EXEC

Command History

| Release | Modification |
|---------------|------------------------------|
| 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.

Use the **show protocols** command to get information about the protocols running on the router and to quickly determine which protocols are active. The command is designed to summarize the important characteristics of the running protocol, and command output varies depending on the specific protocol selected. For BGP, the command output lists the protocol ID, peers with elapsed time since last reset, and miscellaneous information, such as external and internal local distances and sourced routes.

Task ID

| Task ID | Operations |
|---------|------------|
| bgp | read |
| rib | read |

Examples

The following example shows the display for the **show protocols** command using the **bgp** keyword:

```

RP/0/RP0/CPU0:router# show protocols bgp

Routing Protocol "BGP 40"

Address Family IPv4 Unicast:
  Distance: external 20 internal 200 local 200
  Sourced Networks:
    10.100.0.0/16 backdoor
    10.100.1.0/24
    10.100.2.0/24
  Routing Information Sources:
    Neighbor           State/Last update received
    10.5.0.2           Idle
    10.9.0.3           Idle

```

This table describes the significant fields shown in the display.

Table 36: show protocols (BGP) Field Descriptions

| Field | Description |
|-----------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Routing Protocol: | Identifies BGP as the running protocol and displays the BGP AS number. <ul style="list-style-type: none"> • 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. |
| Address Family | Specifies the address family. This can be IPv4 Unicast 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 tcp brief

To display a concise description of TCP connection endpoints, use the **show tcp brief** command in the XR EXEC mode.

show tcp brief [*location node-id*]

| | |
|---------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax Description | location <i>node-id</i> (Optional) Specifies location information for the specified node ID. The node ID variable is mentioned in the <i>rack/slot/module</i> notation. |
|---------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

| | |
|------------------------|-------------------------------|
| Command Default | No default behavior or values |
|------------------------|-------------------------------|

| | |
|----------------------|--------------|
| Command Modes | XR EXEC mode |
|----------------------|--------------|

| | | |
|------------------------|----------------|------------------------------|
| Command History | Release | Modification |
| | Release 3.7.2 | 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 brief** command:

```
RP/0/0/CPU0:ios#show tcp brief
```

| PCB | VRF-ID | Recv-Q | Send-Q | Local Address | Foreign Address | State |
|------------|------------|--------|--------|----------------|-----------------|--------|
| 0x08789b28 | 0x60000000 | 0 | 0 | :::179 | :::0 | LISTEN |
| 0x08786160 | 0x00000000 | 0 | 0 | :::179 | :::0 | LISTEN |
| 0xecb0c9f8 | 0x60000000 | 0 | 0 | 10.0.0.1:12404 | 10.0.0.2:179 | ESTAB |
| 0x0878b168 | 0x60000000 | 0 | 0 | 11.0.0.1:179 | 11.0.0.2:61177 | ESTAB |
| 0xecb0c6b8 | 0x60000000 | 0 | 0 | 0.0.0.0:179 | 0.0.0.0:0 | LISTEN |
| 0x08781590 | 0x00000000 | 0 | 0 | 0.0.0.0:179 | 0.0.0.0:0 | LISTEN |

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 Neighbor group configuration Session group configuration |
|----------------------|---------------------------------------------------------------------------------------|

| | | |
|------------------------|----------------|------------------------------|
| Command History | Release | Modification |
| | 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.

Use the **shutdown** command to terminate any active session for the specified neighbor and remove all associated routing information. Use of the **shutdown** command with a neighbor group or session group may suddenly terminate a large number of BGP neighbor sessions because all neighbors using the neighbor group or session group may be affected.

Use the **show bgp summary** command to display a summary of BGP neighbors. Neighbors that are idle due to the **shutdown** command are displayed with the “Idle (Admin)” state.

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

| | | |
|----------------|----------------|-------------------|
| Task ID | Task ID | Operations |
| | bgp | read, write |

Examples

The following example shows that any active session for neighbor 192.168.40.24 is disabled:

```
RP/0/RP0/CPU0:router(config)# router bgp 1
```

```
RP/0/RP0/CPU0:router(config-bgp)# neighbor 192.168.40.24
RP/0/RP0/CPU0:router(config-bgp-nbr)# shutdown
RP/0/RP0/CPU0:router(config-bgp-nbr)# exit
```

In the following example, the session remains active for neighbor 192.168.40.24 because the inherited **shutdown** command has been overridden:

```
RP/0/RP0/CPU0:router(config)# router bgp 1
RP/0/RP0/CPU0:router(config-bgp)# session-group group1
RP/0/RP0/CPU0:router(config-bgp-sngrp)# shutdown
RP/0/RP0/CPU0:router(config-bgp-sngrp)# exit
RP/0/RP0/CPU0:router(config-bgp)# neighbor 192.168.40.24
RP/0/RP0/CPU0:router(config-bgp-nbr)# remote-as 1
RP/0/RP0/CPU0:router(config-bgp-nbr)# use session-group group1
RP/0/RP0/CPU0:router(config-bgp-nbr)# shutdown inheritance-disable
RP/0/RP0/CPU0:router(config-bgp-nbr)# exit
```

Related Commands

| Command | Description |
|-----------------------------------------------|------------------------------------------------------------------------|
| neighbor-group, on page 141 | Creates a neighbor group and enters neighbor group configuration mode. |
| session-group, on page 189 | Creates a session group and enters session group configuration mode. |
| show bgp summary, on page 320 | Displays the status of all BGP connections. |

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 | <i>socket-size</i> Size (in bytes) of the receive-side socket buffers. Range is 512 to 131072. | |
| | <i>bgp-size</i> (Optional) Size (in bytes) of the receive buffers in BGP. Range is 512 to 131072. | |
| Command Default | <i>socket-size</i> : 32,768 bytes <i>bgp-size</i> : 4,032 bytes | |
| Command Modes | Router configuration | |
| Command History | Release | Modification |
| | 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.

Use the **socket receive-buffer-size** command to increase the buffer size when receiving updates from a neighbor. Using larger buffers can improve convergence time because the software can process more packets simultaneously. However, allocating larger buffers uses more memory on your router.



Note Increasing the socket buffer size uses more memory only when more messages are waiting to be processed by the software. In contrast, increasing the BGP buffer size uses extra memory indefinitely.

Use the **receive-buffer-size** command on individual neighbors to change the values set by the **socket receive-buffer-size** command.

| Task ID | Task ID | Operations |
|---------|---------|----------------|
| | bgp | read, write |

Examples

The following example shows how to set the receive buffer sizes for all neighbors to 65,536 bytes for the socket buffer and 8192 bytes for the BGP buffer:

```
RP/0/RP0/CPU0:router(config)# router bgp 1  
RP/0/RP0/CPU0:router(config-bgp)# socket receive-buffer-size 65536 8192
```

Related Commands

| Command | Description |
|------------------------------------------------------|----------------------------------------------------------|
| receive-buffer-size, on page 167 | Sets the size of the receive buffers for a BGP neighbor. |
| socket send-buffer-size, on page 343 | 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 | <i>socket-size</i> Size (in bytes) of the send-side socket buffers. Range is 4096 to 131072. | | | | |
|---------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------|--------------|---------------|------------------------------|
| | <i>bgp-size</i> (Optional) Size (in bytes) of the send buffers in BGP. Range is 4096 to 131072. | | | | |
| Command Default | <i>socket-size</i> : 10240 bytes <i>bgp-size</i> : 4096 bytes | | | | |
| Command Modes | Router configuration | | | | |
| Command History | <table border="1"> <thead> <tr> <th>Release</th> <th>Modification</th> </tr> </thead> <tbody> <tr> <td>Release 5.0.0</td> <td>This command was introduced.</td> </tr> </tbody> </table> | Release | Modification | Release 5.0.0 | This command was introduced. |
| Release | Modification | | | | |
| 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.

Use the **socket send-buffer-size** command to increase the buffer size when sending updates to neighbors. Using larger buffers can improve convergence time because the software can process more packets simultaneously. However, allocating larger buffers uses more memory on your router.



Note Increasing the socket buffer size uses more memory only when more messages are waiting to be sent by the software. In contrast, increasing the BGP buffer size uses extra memory indefinitely.

Use the **send-buffer-size** command on individual neighbors to change the values set by the **socket send-buffer-size** command.

| Task ID | Task ID | Operations |
|---------|---------|----------------|
| | bgp | read, write |

Examples

The following example shows how to set the send buffer sizes for all neighbors to 8192 bytes for the socket buffer and the BGP buffer:

```
RP/0/RP0/CPU0:router(config)# router bgp 1  
RP/0/RP0/CPU0:router(config-bgp)# socket send-buffer-size 8192 8192
```


Related Commands

| Command | Description |
|---------------------------------------------------------|-------------------------------------------------------------|
| send-buffer-size, on page 183 | Sets the size of the send buffers for a BGP neighbor. |
| socket receive-buffer-size, on page 341 | 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}]
no soft-reconfiguration inbound [{always | inheritance-disable}]
```

| Syntax Description | <p>always (Optional) Always performs a soft inbound clear using stored updates, even if the neighbor supports the route refresh capability.</p> <p>inheritance-disable (Optional) Overrides configuration for this command that may be inherited from a neighbor group or address family group.</p> | | | | |
|-------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------|--------------|---------------|------------------------------|
| Command Default | Soft reconfiguration is not enabled. | | | | |
| Command Modes | <p>IPv4 address family group configuration</p> <p>IPv6 address family group configuration</p> <p>IPv4 neighbor address family configuration</p> <p>IPv4 neighbor group address family configuration</p> <p>IPv6 neighbor group address family configuration</p> | | | | |
| Command History | <table border="1"> <thead> <tr> <th>Release</th> <th>Modification</th> </tr> </thead> <tbody> <tr> <td>Release 5.0.0</td> <td>This command was introduced.</td> </tr> </tbody> </table> | Release | Modification | Release 5.0.0 | This command was introduced. |
| Release | Modification | | | | |
| Release 5.0.0 | This command was introduced. | | | | |
| Usage Guidelines | <p>To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.</p> <p>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.</p> | | | | |
|  Note | If an address family group, neighbor group, or session group is configured, the configuration inside these configuration groups will not be effective unless it is applied directly or indirectly to one or more neighbors. | | | | |



Note The `bgp auto-policy-soft-reset` is enabled by default. A soft clear is done automatically when the inbound policy configured with the `route-policy` (BGP) command is changed. This behavior can be changed by disabling the auto-policy-soft-reset using the `bgp auto-policy-soft-reset disable` command.

If the neighbor supports the route refresh capability, then the original routes are not stored because they can be retrieved from the neighbor through a route refresh request. However, if the `always` keyword is specified, the original routes are stored even when the neighbor supports the route refresh capability.

If the `soft-reconfiguration inbound` command is not configured and the neighbor does not support the route refresh capability, then an inbound soft clear is not possible. In that case, the only way to rerun the inbound policy is to use the `clear bgp ip-address` command to reset the neighbor BGP session.



Note If there is an existing BGP session with a neighbor that does not support the route refresh capability, the session is terminated and a new one is initiated.



Note The extra routes stored as a result of configuring this command use more memory on the router.

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

Task ID

| Task ID | Operations |
|---------|----------------|
| bgp | read, write |

Examples

The following example shows inbound soft reconfiguration enabled for IP Version 4 (IPv4) unicast routes received from neighbor 10.108.1.1. The software stores all routes received in their unmodified form so that when an inbound soft clear is performed later, the stored information can then be used to generate a new set of modified routes.

```
RP/0/RP0/CPU0:router(config)# router bgp 100
RP/0/RP0/CPU0:router(config-bgp)# neighbor 10.108.1.1
RP/0/RP0/CPU0:router(config-bgp-nbr)# remote-as 100
RP/0/RP0/CPU0:router(config-bgp-nbr)# address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-bgp-nbr-af)# soft-reconfiguration inbound
RP/0/RP0/CPU0:router(config-bgp-nbr-af)# exit
```

Related Commands

| Command | Description |
|--------------------------------------|-------------------------------------------------------------------------------------------------------|
| af-group, on page 10 | Creates an address family group for BGP neighbors and enters address family group configuration mode. |

| Command | Description |
|----------------------------------------------------------------|-----------------------------------------------------------------------------------------------------|
| bgp auto-policy-soft-reset disable, on page 29 | Disables an automatic soft reset of BGP peers when the configured inbound route policy is modified. |
| clear bgp, on page 82 | Resets a BGP connection using a soft or hard reset. |
| neighbor-group, on page 141 | Creates a neighbor group and enters neighbor group configuration mode. |
| route-policy (BGP), on page 178 | 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 | <i>id</i> ID of the speaker process. Range is 1 to 15. |
|---------------------------|--------------------------------------------------------|

| | |
|------------------------|---------------|
| Command Default | Default is 0. |
|------------------------|---------------|

| | |
|----------------------|-----------------------------|
| Command Modes | Session group configuration |
|----------------------|-----------------------------|

| Command History | Release | Modification |
|------------------------|----------------|------------------------------|
| | 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 | Operations |
|----------------|----------------|-------------------|
| | bgp | read, write |

| | |
|-----------------|------------------------------------------------------------------------------------------|
| Examples | The following example shows how to allocate speaker process 3 to neighbor 192.168.40.24: |
|-----------------|------------------------------------------------------------------------------------------|

```
RP/0/RP0/CPU0:router(config)# router bgp 109
RP/0/RP0/CPU0:router(config-bgp)# neighbor 192.168.40.24
RP/0/RP0/CPU0:router(config-bgp-nbr)# speaker-id 3
```


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

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 5.0.1 | 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.



Note

Table policy provides users with the ability to drop routes from the RIB based on match criteria. This feature can be useful in certain applications and should be used with caution as it can easily create a routing ‘black hole’ where BGP advertises routes to neighbors that BGP does not install in its global routing table and forwarding table.

Use the **table-policy** command to modify route attributes as the routes are installed into the routing table by Border Gateway Protocol (BGP). Commonly, it is used to set the traffic index attribute.

Task ID

| Task ID | Operations |
|---------|----------------|
| bgp | read, write |

Examples

The following example shows how to apply the set-traffic-index policy to IPv4 unicast routes being installed into the routing table:

```
RP/0/RP0/CPU0:router(config)# router bgp 1
RP/0/RP0/CPU0:router(config-bgp)# address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-bgp-af)# table-policy set-traffic-index
```

Related Commands

| Command | Description |
|---------------------------|--------------------------------------------------------------------|
| route-policy (RPL) | Defines a route policy and enters route policy configuration mode. |

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

Neighbor group configuration

Session group configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| 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.

The timers actually used in connection with the neighbor may not be the same as those configured with this command. The actual timers are negotiated with the neighbor when establishing the session. The negotiated hold time is the minimum of the configured time and the hold time received from the neighbor. If the negotiated hold time is 0, keepalives are disabled.

The configured value for the keepalive must not exceed one-third of the negotiated hold time. If it does, a value of one-third of the negotiated hold time is used.

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

In cases where mechanisms such as Bi-directional Forwarding Detection (BFD), BGP fast-external-failover or Next-hop Tracking cannot be employed to detect and react to changes in the network in a faster manner, BGP Keepalive and Hold-timer values can be configured to use smaller values than the default (60 and 180 seconds respectively). When using aggressive values, consider the router's profile and scale, particularly in respect to the number of BGP neighbours that will be using sessions with the non-default timers.

Sessions using very aggressive values will be more susceptible to flap during events that cause the Route-Processor's CPU utilization levels to increase. Such events include component OIR, Route-Processor Failover, network instability, excessive churn in routing protocols etc. It is therefore recommended that the desired scale and profile of the router be tested with the non-default timer values, subjecting the router to CPU-intensive events in order to determine the timer threshold values that are appropriate for the router before configuring the values in an operational network.

The BGP Non-Stop Routing (NSR) is able to sustain sessions with more aggressive timer values than BGP Graceful Restart (GR) since in the event of a Route-Processor Failover, Graceful Restart (GR) requires the re-establishment of the TCP session over which the BGP session takes place. When using Non-Stop Routing (NSR), both the underlying TCP session and BGP session are maintained during Route-Processor failover.

| Task ID | Task ID | Operations |
|---------|---------|----------------|
| | bgp | read, write |

Examples

The following example shows how to change the keepalive timer to 70 seconds and the hold-time timer to 210 seconds for the BGP peer 192.168.40.24:

```
RP/0/RP0/CPU0:router(config)# router bgp 109
RP/0/RP0/CPU0:router(config-bgp)# neighbor 192.168.40.24
RP/0/RP0/CPU0:router(config-bgp-nbr)# remote-as 1
RP/0/RP0/CPU0:router(config-bgp-nbr)# timers 70 210
```

Related Commands

| Command | Description |
|---------------------------------------------|-------------------------------------------------------------------------------------------------------|
| af-group, on page 10 | Creates an address family group for BGP neighbors and enters address family group configuration mode. |
| neighbor-group, on page 141 | Creates a neighbor group and enters neighbor group configuration mode. |
| session-group, on page 189 | Creates a session group and enters session group configuration mode. |
| timers bgp, on page 353 | 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 | <i>keepalive</i> Frequency (in seconds) with which the software sends keepalive messages to a neighbor. Range is 0 to 65535. | |
| | <i>hold-time</i> 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 | <i>keepalive</i> : 60 seconds <i>hold-time</i> : 180 seconds | |
| Command Modes | Router configuration | |
| Command History | Release | Modification |
| | 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. | |
| | Use the timers bgp command to adjust the default timer times used by all BGP neighbors. The values can be overridden on particular neighbors using the timers command in the neighbor configuration mode. | |
| | The timers actually used in connection with the neighbor may not be the same as those configured with this command. The actual timers are negotiated with the neighbor when establishing the session. The negotiated hold time is the minimum of the configured time and the hold time received from the neighbor. If the negotiated hold time is 0, keepalives are disabled. | |
| | The configured value for the keepalive must not exceed one-third of the negotiated hold time. If it does, a value of one-third of the negotiated hold time is used. | |
| In cases where mechanisms such as Bi-directional Forwarding Detection (BFD), BGP fast-external-failover or Next-hop Tracking cannot be employed to detect and react to changes in the network in a faster manner, BGP Keepalive and Hold-timer values can be configured to use smaller values than the default (60 and 180 seconds respectively). When using aggressive values, consider the router's profile and scale, particularly in respect to the number of BGP neighbors that will be using sessions with the non-default timers. | | |
| Sessions using very aggressive values will be more susceptible to flap during events that cause the Route-Processor's CPU utilization levels to increase. Such events include component OIR, Route-Processor Failover, network instability, excessive churn in routing protocols etc. It is therefore recommended that the desired scale and profile of the router be tested with the non-default timer values, subjecting the router to | | |

CPU-intensive events in order to determine the timer threshold values that are appropriate for the router before configuring the values in an operational network.

The BGP Non-Stop Routing (NSR) is able to sustain sessions with more aggressive timer values than BGP Graceful Restart (GR) since in the event of a Route-Processor Failover, Graceful Restart (GR) requires the re-establishment of the TCP session over which the BGP session takes place. When using Non-Stop Routing (NSR), both the underlying TCP session and BGP session are maintained during Route-Processor failover.

| Task ID | Task ID | Operations |
|---------|---------|----------------|
| | bgp | read, write |

Examples

The following example shows how to configure a default keepalive time of 30 seconds and a default hold time of 90 seconds:

```
RP/0/RP0/CPU0:router(config)# router bgp 1
RP/0/RP0/CPU0:router(config-bgp)# timers bgp 30 90
```

Related Commands

| Command | Description |
|-------------------------------------------|------------------------------------------------|
| timers (BGP), on page 351 | Adjusts BGP network timers for a BGP neighbor. |

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 Neighbor group configuration Session group configuration | | | | |
| Command History | <table border="1"> <thead> <tr> <th>Release</th> <th>Modification</th> </tr> </thead> <tbody> <tr> <td>Release 5.0.0</td> <td>This command was introduced.</td> </tr> </tbody> </table> | Release | Modification | Release 5.0.0 | This command was introduced. |
| Release | Modification | | | | |
| Release 5.0.0 | This command was introduced. | | | | |
| Usage Guidelines | <p>To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.</p> <p>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.</p> <p>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.</p> <p>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.</p> | | | | |



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/RP0/CPU0:router(config)# router bgp 65534
RP/0/RP0/CPU0:router(config-bgp)# neighbor 192.168.223.7
RP/0/RP0/CPU0:router(config-bgp-nbr)# remote-as 65507
RP/0/RP0/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/RP0/CPU0:router(config)# router bgp 65534
RP/0/RP0/CPU0:router(config-bgp)# session-group ebgp-nbrs
RP/0/RP0/CPU0:router(config-bgp-sngrp)# ttl-security
RP/0/RP0/CPU0:router(config-bgp-sngrp)# exit
RP/0/RP0/CPU0:router(config-bgp)# neighbor 192.168.223.1
RP/0/RP0/CPU0:router(config-bgp-nbr)# remote-as 65501
RP/0/RP0/CPU0:router(config-bgp-nbr)# use session-group ebgp-nbrs
RP/0/RP0/CPU0:router(config-bgp-nbr)# exit
RP/0/RP0/CPU0:router(config-bgp)# neighbor 192.168.223.2
RP/0/RP0/CPU0:router(config-bgp-nbr)# remote-as 65502
RP/0/RP0/CPU0:router(config-bgp-nbr)# use session-group ebgp-nbrs
RP/0/RP0/CPU0:router(config-bgp-nbr)# exit
RP/0/RP0/CPU0:router(config-bgp)# neighbor 192.168.223.3
```



```
RP/0/RP0/CPU0:router(config-bgp-nbr)# remote-as 65503
RP/0/RP0/CPU0:router(config-bgp-nbr)# use session-group ebgp-nbrs
RP/0/RP0/CPU0:router(config-bgp-nbr)# exit
```

Related Commands

| Command | Description |
|---------------------------------------------|--------------------------------------------------------------------------------------------------------------|
| ebgp-multihop, on page 116 | Accepts and attempts BGP connections to external peers residing on networks that are not directly connected. |
| neighbor-group, on page 141 | Creates a neighbor group and enters neighbor group configuration mode. |
| session-group, on page 189 | 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 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 | <i>update-limit-MB</i> 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 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. |
|-------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

Use the **update limit** command to configure a global limit on the size of messages the software queues when updating peers. Increasing the limit can result in faster Border Gateway Protocol (BGP) convergence, but also may result in higher memory usage during convergence.

| | | |
|----------------|----------------|------------------|
| Task ID | Task ID | Operation |
| | bgp | read, write |

This example shows how to set the update limit as 1024 MB:

```
RP/0/RP0/CPU0:router(config)# router bgp 65000
RP/0/RP0/CPU0:router(config-bgp)#update limit 1024
```

| | | |
|-------------------------|----------------------------------------------------------|-----------------------------------------------------------------------------------------------|
| Related Commands | Command | Description |
| | update limit address-family, on page 359 | Sets upper bound on transient memory usage for update generation for an address family. |
| | update limit sub-group, on page 360 | Sets upper bound on transient memory usage for update generation for eBGP or iBGP sub-groups. |

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 | <i>update-limit-MB</i> 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 | |
| Command History | Release | Modification |
| | Release 5.0.0 | This command was introduced. |
| Usage Guidelines | <p>To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.</p> <p>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.</p> | |
| Task ID | Task ID | Operation |
| | bgp | read, write |
| | <p>This example shows how to set the update limit as <i>512 MB</i> for address family IPv4 unicast:</p> <pre>RP/0/RP0/CPU0:router#configure RP/0/RP0/CPU0:router(config)#router bgp 100 RP/0/RP0/CPU0:router(config-bgp)#address-family ipv4 unicast RP/0/RP0/CPU0:router(config-bgp-af)#update limit address-family 512</pre> | |
| Related Commands | Command | Description |
| | update limit, on page 358 | Sets upper bound on transient memory usage for update generation. |
| | update limit sub-group, on page 360 | Sets upper bound on transient memory usage for update generation for eBGP or iBGP sub-groups. |

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 | Parameter | Description |
|--------------------|------------------------|-------------------------------------------------------------------|
| | ebgp | Specifies the update limit for eBGP sub-groups. |
| | ibgp | Specifies the update limit for iBGP sub-groups. |
| | <i>update-limit-MB</i> | 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

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | 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.

Use the **update limit sub-group** command to configure a global limit on the size of messages the software queues when updating peers. Increasing the limit can result in faster Border Gateway Protocol (BGP) convergence, but also may result in higher memory usage during convergence.

| Task ID | Task ID | Operation |
|---------|---------|----------------|
| | bgp | read, write |

This example shows how to set the update limit as *256 MB* for eBGP sub-group under address family IPV4 unicast:

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

```
RP/0/RP0/CPU0:router(config-bgp-af)#update limit sub-group ebgp 256
```

Related Commands

| Command | Description |
|----------------------------------------------------------|-----------------------------------------------------------------------------------------|
| update limit, on page 358 | Sets upper bound on transient memory usage for update generation. |
| update limit address-family, on page 359 | Sets upper bound on transient memory usage for update generation for an address family. |

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 | <table border="1"> <thead> <tr> <th>Release</th> <th>Modification</th> </tr> </thead> <tbody> <tr> <td>Release 5.0.0</td> <td>This command was introduced.</td> </tr> </tbody> </table> | Release | Modification | Release 5.0.0 | This command was introduced. |
| Release | Modification | | | | |
| 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 | <table border="1"> <thead> <tr> <th>Task ID</th> <th>Operation</th> </tr> </thead> <tbody> <tr> <td>bgp</td> <td>read, write</td> </tr> </tbody> </table> | Task ID | Operation | bgp | read, write |
| Task ID | Operation | | | | |
| bgp | read, write | | | | |

This example shows how to disable inbound update message basic error handling for eBGP neighbors:

```
RP/0/RP0/CPU0:router#configure
RP/0/RP0/CPU0:router (config)#router bgp 100
RP/0/RP0/CPU0:router (config-bgp)#update in error-handling basic ebgp disable
```

This example shows how to disable inbound update message basic error handling for iBGP neighbors:

```
RP/0/RP0/CPU0:router#configure
RP/0/RP0/CPU0:router (config)#router bgp 100
RP/0/RP0/CPU0:router (config-bgp)#update in error-handling basic ibgp disable
```

update in error-handling extended

To enable inbound update message extended error handling for eBGP or iBGP neighbors, use the **update in error-handling extended** command in 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 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 |
| | bgp | read, write |

This example shows how to enable inbound update message extended error handling for eBGP neighbors:

```
RP/0/RP0/CPU0:router#configure
RP/0/RP0/CPU0:router(config)#router bgp 100
RP/0/RP0/CPU0:router(config-bgp)#update in error-handling extended ebgp
```

This example shows how to enable inbound update message extended error handling for iBGP neighbors:

```
RP/0/RP0/CPU0:router#configure
RP/0/RP0/CPU0:router(config)#router bgp 100
RP/0/RP0/CPU0:router(config-bgp)#update in error-handling extended ibgp
```

update in labeled-unicast equivalent

To treat SAFIs unicast and labelled unicast for a given neighbor as same, use the **update in labeled-unicast equivalent** command in neighbor configuration mode. This command was introduced for interoperability between Cisco IOS XR and IOS XE systems. For example, If the XE router advertises ipv4 unicast and then withdraw with the ipv4 label-unicast of the same route, with knob enabled, then the BGP in the receiving XR router treats both of them same. Hence BGP removes the prefix information from its database.

update in labeled-unicast equivalent [inheritance-disable]

| | | |
|---------------------------|----------------------------|-------------------------------------------------------------------------------------------------|
| Syntax Description | equivalent | IPv4/IPv6 labeled-unicast inbound updates (paths) treated equivalent to unicast updates (paths) |
| | inheritance-disable | Does not inherit this unicast configuration from parent group |

Command Default None

Command Modes Neighbor configuration mode
Neighbor group configuration mode

| Command History | Release | Modification |
|------------------------|----------------|------------------------------|
| | Release 7.0.1 | This command was introduced. |

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

| Task ID | Task ID | Operation |
|----------------|----------------|------------------|
| | bgp | read, write |

Example

The following example shows how to treat IPv4 labeled-unicast inbound updates equivalent to unicast updates (paths):

```
RP/0/RP0/CPU0:router#configure
RP/0/RP0/CPU0:router (config)#router bgp 100
RP/0/RP0/CPU0:router (config-bgp)#neighbor 10.2.3.4
RP/0/RP0/CPU0:router (config-bgp-nbr)#update in labeled-unicast equivalentexit
```


update out logging

To enable logging of update generation events, use the **update out logging** command in 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 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 |
|---------|---------|----------------|
| | bgp | read, write |

This example shows how to enable logging of update generation events:

```
RP/0/RP0/CPU0:router#configure
RP/0/RP0/CPU0:router(config)#router bgp 100
RP/0/RP0/CPU0:router(config-bgp)#update out logging
```

update-source

To allow internal Border Gateway Protocol (iBGP) sessions to use the primary IP address from a particular interface as the local address when forming an iBGP session with a neighbor, use the **update-source** command in an appropriate configuration mode. To set the chosen local IP address to the nearest interface to the neighbor, use the **no** form of this command.

update-source *type interface-path-id*
no update-source [*type interface-path-id*]

Syntax Description

type Interface type. For more information, use the question mark (?) online help function.

interface-path-id Physical interface or virtual interface.

Note Use the **show interfaces** command to see a list of all interfaces currently configured on the router.

For more information about the syntax for the router, use the question mark (?) online help function.

Command Default

Best local address

Command Modes

Neighbor configuration
 Neighbor group configuration
 Session group configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| 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.

The **update-source** command is commonly used with the loopback interface feature for iBGP sessions. The loopback interface is defined, and the interface address is used as the endpoint for a BGP session through the **update-source** command. This mechanism allows a BGP session to remain up even if the outbound interface goes down, provided there is another route to the neighbor.

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

Task ID

| Task ID | Operations |
|---------|----------------|
| bgp | read, write |

Examples

The following example shows how to configure this router to use the IP address from the Loopback0 interface when trying to open a session with neighbor 172.20.16.6:

```
RP/0/RP0/CPU0:router(config)# router bgp 110  
RP/0/RP0/CPU0:router(config-bgp)# neighbor 172.20.16.6  
RP/0/RP0/CPU0:router(config-bgp-nbr)# remote-as 110  
RP/0/RP0/CPU0:router(config-bgp-nbr)# update-source Loopback0
```

Related Commands

| Command | Description |
|---------------------------------------------|------------------------------------------------------------------------|
| neighbor-group, on page 141 | Creates a neighbor group and enters neighbor group configuration mode. |
| session-group, on page 189 | Creates a session group and enters session group configuration mode. |

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-name }
no use {af-group [group-name] | neighbor-group [group-name] | session-group [group-name] }
```

| Syntax Description | |
|-----------------------|------------------------------------------------------------------------------------------------------------------|
| af-group | Specifies an address family group. |
| <i>group-name</i> | 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

For **use session-group** version:

- Neighbor group configuration
- Neighbor configuration
- Session-group configuration

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | 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.

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.

```
RP/0/RP0/CPU0:router (config) # router bgp 1
RP/0/RP0/CPU0:router (config-bgp) # session-group session1
RP/0/RP0/CPU0:router (config-bgp-sngrp) # advertisement-interval 40
RP/0/RP0/CPU0:router (config-bgp-sngrp) # timers 30 90
RP/0/RP0/CPU0:router (config-bgp-sngrp) # exit
RP/0/RP0/CPU0:router (config-bgp) # neighbor 172.168.40.24
RP/0/RP0/CPU0:router (config-bgp-nbr) # remote-as 2
RP/0/RP0/CPU0:router (config-bgp-nbr) # use session-group session1
RP/0/RP0/CPU0:router (config-bgp-nbr) # exit
```

The following example is similar to the previous example, but in this case the **timers** command on the session group does not take effect on the neighbor because it is overridden by a **timers** command directly configured for the neighbor.

```
RP/0/RP0/CPU0:router (config) # router bgp 1
RP/0/RP0/CPU0:router (config-bgp) # session-group session1
RP/0/RP0/CPU0:router (config-bgp-sngrp) # advertisement-interval 40
RP/0/RP0/CPU0:router (config-bgp-sngrp) # timers 30 90
RP/0/RP0/CPU0:router (config-bgp-sngrp) # exit
RP/0/RP0/CPU0:router (config-bgp) # neighbor 172.168.40.24
RP/0/RP0/CPU0:router (config-bgp-nbr) # remote-as 2
RP/0/RP0/CPU0:router (config-bgp-nbr) # use session-group session1
RP/0/RP0/CPU0:router (config-bgp-nbr) # timers 60 180
RP/0/RP0/CPU0:router (config-bgp-nbr) # exit
```

The following example shows a neighbor inheriting configuration from a session group that likewise inherits configuration from another session group. The configuration from both session groups take effect on the neighbor:

```
RP/0/RP0/CPU0:router (config) # router bgp 1
RP/0/RP0/CPU0:router (config-bgp) # session-group session1
RP/0/RP0/CPU0:router (config-bgp-sngrp) # advertisement-interval 40
RP/0/RP0/CPU0:router (config-bgp-sngrp) # exit
RP/0/RP0/CPU0:router (config-bgp) # session-group session2
RP/0/RP0/CPU0:router (config-bgp-sngrp) # use session-group session1
RP/0/RP0/CPU0:router (config-bgp-sngrp) # update-source Loopback0
RP/0/RP0/CPU0:router (config-bgp-sngrp) # exit
RP/0/RP0/CPU0:router (config-bgp) # neighbor 172.168.40.24
RP/0/RP0/CPU0:router (config-bgp-nbr) # remote-as 1
RP/0/RP0/CPU0:router (config-bgp-nbr) # use session-group session2
RP/0/RP0/CPU0:router (config-bgp-nbr) # exit
```

| Related Commands | Command | Description |
|------------------|------------------------------------------------------|-------------------------------------------------------------------------------------------------------|
| | af-group, on page 10 | Creates an address family group for BGP neighbors and enters address family group configuration mode. |
| | session-group, on page 189 | Creates a session group and enters session group configuration mode. |
| | neighbor-group, on page 141 | Creates a neighbor group and enters neighbor group configuration mode. |
| | remote-as (BGP), on page 173 | Creates a BGP neighbor and begins the exchange of routing information. |
| | show bgp af-group, on page 220 | Displays information about BGP configuration for address family groups. |
| | show bgp neighbor-group, on page 252 | Displays information about the BGP configuration for neighbor groups. |
| | show bgp neighbors, on page 256 | Displays information about BGP neighbors. |
| | show bgp session-group, on page 316 | Displays information about the BGP configuration for session groups. |

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

Command History

Release

Release 5.0.0

Modification

This command was introduced.

Usage Guidelines

To use this command, you must be in a user group associated with a task 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* module of the *Routing Configuration Guide for Cisco NCS 6000 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 101](#) command.

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

Task ID

| Task ID | Operations |
|---------|----------------|
| bgp | read, write |

Examples

The following example shows how to assign a weight of 50 to all IP Version 4 (IPv4) unicast routes learned through 172.20.16.6:

```
RP/0/RP0/CPU0:router(config)# router bgp 1
RP/0/RP0/CPU0:router(config-bgp)# neighbor 172.20.16.6
RP/0/RP0/CPU0:router(config-bgp-nbr)# remote-as 1
RP/0/RP0/CPU0:router(config-bgp-nbr)# address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-bgp-nbr-af)# weight 50
RP/0/RP0/CPU0:router(config-bgp-nbr-af)# exit
```

Related Commands

| Command | Description |
|---------------------------------------------|-------------------------------------------------------------------------------------------------------|
| af-group, on page 10 | Creates an address family group for BGP neighbors and enters address family group configuration mode. |
| clear bgp, on page 82 | Resets a group of BGP neighbors. |
| neighbor-group, on page 141 | Creates a neighbor group and enters neighbor group configuration mode. |
| session-group, on page 189 | 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 NCS 6000 Series Router.

- [class-map type traffic \(BGP-flowspec\), on page 376](#)
- [class type traffic, on page 377](#)
- [drop \(BGP-flowspec\), on page 378](#)
- [flowspec, on page 379](#)
- [flowspec disable, on page 380](#)
- [local-install, on page 381](#)
- [match destination-address, on page 382](#)
- [match destination-port, on page 383](#)
- [match dscp, on page 384](#)
- [match fragment-type, on page 387](#)
- [match icmp code, on page 388](#)
- [match icmp type, on page 389](#)
- [match packet length, on page 390](#)
- [match protocol, on page 391](#)
- [match source-address, on page 393](#)
- [match source-port, on page 394](#)
- [match tcp flag, on page 395](#)
- [policy-map, on page 396](#)
- [redirect \(BGP Flowspec\), on page 398](#)
- [service-policy, on page 399](#)
- [show flowspec, on page 400](#)

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 in XR Config 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. |
|--------------------|-----------------------|-------------------------------------------------|
| | <i>class-map-name</i> | Name of the class for the class map. |

Command Default None

Command Modes XR Config

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | Release 5.2.4 | 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.

This example shows how to specify class305 as the name of a class and defines a class map for this class.

```
RP/0/RP0/CPU0:router# config
RP/0/RP0/CPU0:router(config)# class-map type traffic match-all class305
RP/0/RP0/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 | <i>class-name</i> 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 | <table border="1"> <thead> <tr> <th>Release</th> <th>Modification</th> </tr> </thead> <tbody> <tr> <td>Release 5.2.0</td> <td>This command was introduced.</td> </tr> </tbody> </table> | Release | Modification | Release 5.2.0 | This command was introduced. |
| Release | Modification | | | | |
| Release 5.2.0 | This command was introduced. | | | | |
| Usage Guidelines | <p>To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.</p> <p>This example shows how to associate a class map with the policy map:</p> <pre>RP/0/RP0/CPU0:router# config RP/0/RP0/CPU0:router(config)# policy-map type pbr p1 RP/0/RP0/CPU0:router(config-pmap)# class type traffic c1 RP/0/RP0/CPU0:router(config-pmap-c)# set dscp 34</pre> | | | | |

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 5.2.4 | 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.

Examples

This example shows how to discard packets:

```
RP/0/RP0/CPU0:router#config
RP/0/RP0/CPU0:router(config)# policy-map type pbr match_dest_110.1.1.x_drop
RP/0/RP0/CPU0:router(config-pmap)# class type traffic match_dest_110.1.1.x
RP/0/RP0/CPU0:router(config-pmap-c)# drop
```

flowspec

To enter BGP flowspec configuration mode, use the **flowspec** command in XR Config mode.

flowspec

Syntax Description This command has no keywords or arguments.

Command Default No default behavior or values

Command Modes XR Config

| Command History | Release | Modification |
|------------------------|----------------|------------------------------|
| | Release 5.2.4 | 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.

Examples

This example show how to enter flowspec configuration mode.

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# flowspec
RP/0/RP0/CPU0:router(config-flowspec)#
```

flowspec disable

To disable flowspec configuration on all interfaces, use the **flowspec disable** command in interface configuration mode.

ipv4
flowspec disable

| | | |
|---------------------------|-------------------------------|------------------------------|
| Syntax Description | ipv4 | Specifies IPv4 interfaces. |
| Command Default | No default behavior or values | |
| Command Modes | Interface configuration | |
| Command History | Release | Modification |
| | Release 5.2.4 | 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.

Examples

This example shows how to disable flowspec configuration on all interfaces.

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# interface GigabitEthernet 0/2/0/2
RP/0/RP0/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
 VRF IPv4 address family configuration

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | Release 5.2.4 | 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.

Examples This example show how to install flowspec policy on all interfaces under flowspec subaddress family configuration mode.

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# flowspec
RP/0/RP0/CPU0:router(config-flowspec)# address-family ipv4
RP/0/RP0/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} address
no match destination-address {ipv4} address
```

| | | |
|---------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------|
| Syntax Description | <p>ipv4 Indicates an IPv4 address.</p> <p>address Specifies a destination address.</p> | |
| Command Default | No default behavior or values | |
| Command Modes | Class map configuration | |
| Command History | Release | Modification |
| | Release 5.2.4 | 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. | |
| Examples | This example shows how to match a destination ipv4 address: | |
| | <pre>RP/0/RP0/CPU0:router(config)#class-map type traffic match-all RP/0/RP0/CPU0:router(config-cmap)# match destination-address ipv4 59.2.1.2 255.255.255.0</pre> | |

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 | [min-value - max-value]}
no match destination-port {destination-port-value | [min-value - max-value]}
```

| Syntax Description | |
|-------------------------------|----------------------------------------------------------------------------|
| <i>destination-port-value</i> | A port Number. Range is from 0 to 65535. |
| <i>min-value</i> | Lower limit of destination port range to match. Value range is 0 to 65535. |
| <i>max-value</i> | Upper limit of destination port range to match. Value range is 0 to 65535. |

Command Default No default behavior or values

Command Modes Class map configuration

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | Release 5.2.4 | 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.

Examples This example shows how to match a destination port:

```
RP/0/RP0/CPU0:router(config)# class-map type traffic match-all
RP/0/RP0/CPU0:router(config-cmap)# match destination-port 1
```

match dscp

To identify specific IP differentiated services code point (DSCP) values as match criteria for a class map, use the **match dscp** command in class map configuration mode. To remove a DSCP value from a class map, use the **no** form of this command.

```
match dscp [{ipv4}] dscp-value [dscp-value1 . . . dscp-value7] |[min-value - max-value]}
no match dscp [{ipv4}] dscp-value [dscp-value1 . . . dscp-value7] |[min-value - max-value]}
```

| Syntax Description | |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| not | (Optional) Negates the specified match result. |
| ipv4 | (Optional) Specifies the IPv4 DSCP value. |
| <i>dscp-value</i> | IP DSCP value identifier that specifies the exact value or a range of values. Range is 0 - 63. Up to eight IP DSCP values can be specified to match packets. Reserved keywords can be specified instead of numeric values. Table 37: IP DSCP Reserved Keywords, on page 385 describes the reserved keywords. |
| <i>min-value</i> | Lower limit of DSCP range to match. Value range is 0 - 63. |
| <i>max-value</i> | Upper limit of DSCP range to match. Value range is 0 - 63. |

Command Default Matching on IP Version 4 (IPv4) packets is the default.

Command Modes Class map configuration

| Command History | Release | Modification |
|-----------------|---------------|-----------------------------------------------------------------|
| | Release 5.0.0 | This command was introduced. |
| | Release 5.2.4 | The <i>min-value</i> and <i>max-value</i> variables were added. |

Usage Guidelines The **match dscp** command specifies a DSCP value that is used as the match criteria against which packets are checked to determine if they belong to the class specified by the class map.

To use the **match dscp** 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 dscp** command in a class map, only the last command entered applies.

The **match dscp** command examines the higher-order six bits in the type of service (ToS) byte of the IP header. Only one of the eight values is needed to yield a match (OR operation).

The command supports only eight IP DSCP values. If you try to configure more match statements after all the eight values are matched, the statements get rejected.

The IP DSCP value is used as a matching criterion only. The value has no mathematical significance. For instance, the IP DSCP value 2 is not greater than 1. The value simply indicates that a packet marked with the IP DSCP value of 2 should be treated differently than a packet marked with an IP DSCP value of 1. The treatment of these marked packets is defined by the user through the setting of policies in policy map class configuration mode.

Table 37: IP DSCP Reserved Keywords

| DSCP Value | Reserved Keyword |
|------------|------------------|
| 0 | default |
| 10 | AF11 |
| 12 | AF12 |
| 14 | AF13 |
| 18 | AF21 |
| 20 | AF22 |
| 22 | AF23 |
| 26 | AF31 |
| 28 | AF32 |
| 30 | AF33 |
| 34 | AF41 |
| 36 | AF42 |
| 38 | AF43 |
| 46 | EF |
| 8 | CS1 |
| 16 | CS2 |
| 24 | CS3 |
| 32 | CS4 |
| 40 | CS5 |
| 48 | CS6 |
| 56 | CS7 |
| ipv4 | ipv4 dscp |

Task ID

Task ID Operations

qos read,
write

Examples

This example shows how to configure the service policy called policy1 and attach service policy policy1 to an interface. In this example, class map dscp14 evaluates all packets entering HundredGigE 0/7/0/0 for an IP DSCP value of 14. If the incoming packet has been marked with the IP DSCP value of 14, the packet is queued to the class queue with the bandwidth setting of 300 kbps.

```
RP/0/RP0/CPU0:router(config)# class-map dscp14
RP/0/RP0/CPU0:router(config-cmap)# match dscp ipv4 14
RP/0/RP0/CPU0:router(config-cmap)# exit

RP/0/RP0/CPU0:router(config)# policy-map policy1
RP/0/RP0/CPU0:router(config-pmap)# class dscp14
RP/0/RP0/CPU0:router(config-pmap-c)# bandwidth 300
RP/0/RP0/CPU0:router(config-pmap-c)# exit
RP/0/RP0/CPU0:router(config-pmap)# exit

RP/0/RP0/CPU0:router(config)# interface HundredGigE 0/7/0/0
RP/0/RP0/CPU0:router(config-if)# service-policy input policy1
```

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 [is-fragment]
no match fragment type [is-fragment]
```

| Syntax Description | is-fragment Matches is-fragment bit. | | | | |
|---------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------|--------------|---------------|------------------------------|
| Command Default | No default behavior or values | | | | |
| Command Modes | Class map configuration | | | | |
| Command History | <table border="1"> <thead> <tr> <th>Release</th> <th>Modification</th> </tr> </thead> <tbody> <tr> <td>Release 5.2.4</td> <td>This command was introduced.</td> </tr> </tbody> </table> | Release | Modification | Release 5.2.4 | This command was introduced. |
| Release | Modification | | | | |
| Release 5.2.4 | 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. | | | | |

Examples

This example shows how to match a fragment-type:

```
RP/0/RP0/CPU0:router(config)# class-map type traffic match-all
RP/0/RP0/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} icmp-code {value | [min-value - max-value]}
no match {ipv4} icmp-code {value | [min-value - max-value]}
```

| Syntax Description | |
|--------------------|-------------------------------------------------------------------|
| ipv4 | Indicates an IPv4 ICMP code. |
| <i>min-value</i> | Lower limit of ICMP type range to match. Value range is 0 to 255. |
| <i>max-value</i> | Upper limit of ICMP type range to match. Value range is 0 to 255. |

Command Default No default behavior or values

Command Modes Class map configuration

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | Release 5.2.4 | 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.

Examples This example shows how to match an IPv4 ICMP code:

```
RP/0/RP0/CPU0:router(config)# class-map type traffic match-all
RP/0/RP0/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} icmp-type {value | [min-value - max-value]}
no match {ipv4} icmp-type {value | [min-value - max-value]}
```

Syntax Description

| | |
|------------------|-------------------------------------------------------------------|
| ipv4 | Indicates an IPv4 ICMP type. |
| <i>min-value</i> | Lower limit of ICMP type range to match. Value range is 0 to 255. |
| <i>max-value</i> | Upper limit of ICMP type range to match. Value range is 0 to 255. |

Command Default

No default behavior or values

Command Modes

Class map configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 5.2.4 | 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.

Examples

This example shows how to match an IPv4 ICMP type:

```
RP/0/RP0/CPU0:router(config)# class-map type traffic match-all
RP/0/RP0/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 | [min-value - max-value]}
no match packet length {value | [min-value - max-value]}
```

| Syntax Description | |
|--------------------|-----------------------------------------------------------|
| <i>value</i> | IP packet length. Range is from 0 to 65535. |
| <i>min-value</i> | Minimum length value to match. Value range is 0 to 65535. |
| <i>max-value</i> | Minimum length value to match. Value range is 0 to 65535. |

Command Default No default behavior or values.

Command Modes Class map configuration

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | Release 5.2.4 | 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.

Examples This example shows how to match a packet length value:

```
RP/0/RP0/CPU0:router(config)# class-map type traffic match-all
RP/0/RP0/CPU0:router(config-cmap)# match packet length 3
```

match protocol

To identify a specific protocol as the match criterion for a class map, use the **match protocol** command in class map configuration mode. To remove protocol-based match criteria from a class map, use the **no** form of this command.

match [**not**] **protocol** {*protocol-value* [*protocol-value1* . . . *protocol-value7*] | [*min-value* - *max-value*]}

no match [**not**] **protocol** {*protocol-value* [*protocol-value1* . . . *protocol-value7*] | [*min-value* - *max-value*]}

| Syntax Description | |
|-----------------------|-----------------------------------------------------------------------------------------------------------------------------------------------|
| not | (Optional) Negates the specified match result. |
| <i>protocol-value</i> | A protocol identifier. A single value for <i>protocol-value</i> (any combination of numbers and names) can be matched in one match statement. |
| <i>min-value</i> | Lower limit of protocol range to match. Value range is 0 - 255. |
| <i>max-value</i> | Upper limit of protocol range to match. Value range is 0 - 255. |

Command Default No default behavior or values

Command Modes Class map configuration

| Command History | Release | Modification |
|-----------------|---------------|-----------------------------------------------------------------|
| | Release 5.0.0 | This command was introduced. |
| | Release 5.2.4 | The <i>min-value</i> and <i>max-value</i> variables were added. |

Usage Guidelines Definitions of traffic classes are based on match criteria, including protocols, access control lists (ACLs), input interfaces, QoS labels, and experimental (EXP) field values. Packets satisfying the match criteria for a class constitute the traffic for that class.

The **match protocol** command specifies the name of a protocol to be used as the match criteria against which packets are checked to determine if they belong to the class specified by the class map. Available protocol names are listed in the table that follows.

The *protocol-value* argument supports a range of protocol numbers. After you identify the class, you may use the **match protocol** command to configure its match criteria.

Table 38: Protocol Names and Descriptions

| Name | Description |
|-------|--------------------------------------------------|
| ahp | Authentication Header Protocol |
| eigrp | Cisco Enhanced Interior Gateway Routing Protocol |

| Name | Description |
|--------|-----------------------------------------------|
| esp | Encapsulation Security Payload |
| gre | Cisco Generic Routing Encapsulation Tunneling |
| icmp | Internet Control Message Protocol |
| igmp | Internet Gateway Message Protocol |
| igrp | Cisco IGRP Routing protocol |
| ipinip | IP in IP tunneling |
| ipv4 | Any IPv4 protocol |
| ipv6 | Any IPv6 protocol |
| mpls | Any MPLS packet |
| nos | KA9Q NOS Compatible IP over IP Tunneling |
| ospf | Open Shortest Path First, Routing Protocol |
| pcp | Payload Compression Protocol |
| pim | Protocol Independent Multicast |
| sctp | Stream Control Transmission Protocol |
| tcp | Transport Control Protocol |
| udp | User Datagram Protocol |

Task ID**Task ID** **Operations**

| | |
|-----|----------------|
| qos | read, write |
|-----|----------------|

Examples

In this example, all TCP packets belong to class class1:

```
RP/0/RP0/CPU0:router(config)# class-map class1
RP/0/RP0/CPU0:router(config-cmap)# match protocol tcp
```

match source-address

To identify a specific source 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 IP address from the matching criteria for a class map, use the **no** form of this command.

```
match source-address {ipv4} address
no match source-address {ipv4} address
```

| Syntax Description | <p>ipv4 Indicates an IPv4 address.</p> <p><i>address</i> Specifies a source address.</p> | | | | |
|---------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------|--------------|---------------|------------------------------|
| Command Default | No default behavior or values | | | | |
| Command Modes | Class map configuration | | | | |
| Command History | <table border="1"> <thead> <tr> <th>Release</th> <th>Modification</th> </tr> </thead> <tbody> <tr> <td>Release 5.2.4</td> <td>This command was introduced.</td> </tr> </tbody> </table> | Release | Modification | Release 5.2.4 | This command was introduced. |
| Release | Modification | | | | |
| Release 5.2.4 | 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. | | | | |
| Examples | <p>This example shows how to match a source ipv4 address:</p> <pre>RP/0/RP0/CPU0:router(config)#class-map type traffic match-all A RP/0/RP0/CPU0:router(config-cmap)# match source-address ipv4 59.2.1.2 255.255.255.0</pre> | | | | |

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 | [min-value - max-value]}
```

```
no match source-port {source-port-value | [min-value - max-value]}
```

Syntax Description

| | |
|--------------------------|-----------------------------------------------------------------------|
| <i>source-port-value</i> | A port Number. Range is from 0 to 65535. |
| <i>min-value</i> | Lower limit of source port range to match. Value range is 0 to 65535. |
| <i>max-value</i> | Upper limit of source port range to match. Value range is 0 to 65535. |

Command Default

No default behavior or values

Command Modes

Class map configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 5.2.4 | 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.

Examples

This example shows how to match a source port:

```
RP/0/RP0/CPU0:router(config)# class-map type traffic match-all
RP/0/RP0/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 valueany
```

| | | |
|---------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------|
| Syntax Description | <p><i>value</i> TCP flag value. Range is from 1 to 4095 (hexadecimal).</p> <p>any Specifies a match based on any bit in the TCP flag.</p> | |
| Command Default | No default behavior or values | |
| Command Modes | Class map configuration | |
| Command History | Release | Modification |
| | Release 5.2.4 | 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. | |
| Examples | This example shows how to match a TCP flag: | |
| | <pre>RP/0/RP0/CPU0:router(config)# class-map type traffic match-all RP/0/RP0/CPU0:router(config-cmap)# match tcp flag 2 any</pre> | |

policy-map

To create or modify a policy map that can be attached to one or more interfaces to specify a service policy, use the **policy-map** command in XR Config mode. To delete a policy map, use the **no** form of this command.

```
policy-map [type qos] policy-name
no policy-map [type qos] policy-name
```

| Syntax Description | | |
|--------------------|--------------------|---------------------------------------------------------------|
| | type qos | (Optional) Specifies type of the service policy. |
| | qos | (Optional) Specifies a quality-of-service (QoS) policy map. |
| | pbr | (Optional) Specifies a policy-based routing (PBR) policy map. |
| | <i>policy-name</i> | Name of the policy map. |

Command Default A policy map does not exist until one is configured. Because a policy map is applied to an interface, no restrictions on the flow of data are applied to any interface until a policy map is created.

Type is QoS when not specified.

Command Modes XR Config mode

| Command History | Release | Modification |
|-----------------|---------------|-----------------------------------|
| | Release 5.0.0 | This command was introduced. |
| | Release 5.2.4 | The pbr keyword was added. |

Usage Guidelines Use the **policy-map** command to specify the name of the policy map to be created, added to, or modified before you can configure policies for classes whose match criteria are defined in a class map. Entering the **policy-map** command enables policy map configuration mode in which you can configure or modify the class policies for that policy map.

You can configure class policies in a policy map only if the classes have match criteria defined for them. Use the **class-map** and **match** commands to configure the match criteria for a class. Because you can configure a maximum of 1024 classes in one policy map, no policy map can contain more than 1024 class policies. The maximum number of 1024 classes per policy includes the implicit default class and its child policies.

A single policy map can be attached to multiple interfaces concurrently.

| Task ID | Task ID | Operations |
|---------|---------|----------------|
| | qos | read, write |

Examples

These examples show how to create a policy map called policy1 and configures two class policies included in that policy map. The policy map is defined to contain policy specification for class1 and the default class (called class-default) to which packets that do not satisfy configured match criteria are directed. Class1 specifies policy for traffic that matches access control list 136.

```
RP/0/RP0/CPU0:router(config)# class-map class1
RP/0/RP0/CPU0:router(config-cmap)# match access-group ipv4 136

RP/0/RP0/CPU0:router(config)# policy-map policy1
RP/0/RP0/CPU0:router(config-pmap)# class class1
RP/0/RP0/CPU0:router(config-pmap-c)# police rate 250

RP/0/RP0/CPU0:router(config-pmap-c)# exit

RP/0/RP0/CPU0:router(config-pmap)# class class-default
RP/0/RP0/CPU0:router(config-pmap-c)# queue-limit 1000000 bytes
```

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 | 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 |
| <i>IPv4 address</i> | | Input IPv4 Nexthop address |
| route-target | | Enter specific route-target string |
| <i>AS-number: index</i> | | Enter 2-byte or 4-byte autonomous system number (AS) and <i>index</i> in hexa decimal or decimal format. |
| <i>IPv4-address: index</i> | | Enter IPv4 address and <i>index</i> in hexa decimal or decimal format. |
| vrf <i>vrf-name</i> | | Enter specific VRF name for the nexthop. |

Command Default None

Command Modes Policy-map configuration

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | Release 5.2.4 | 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/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# policy-map type pbr test1
RP/0/RP0/CPU0:router(config-pmap)# class type traffic test1
RP/0/RP0/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. |
| | <i>policy-name</i> | Name of the policy map. |

Command Default No default behavior or values

Command Modes IPv4 address family configuration
VRF IPv4 address family configuration

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | Release 5.2.4 | 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.

Examples

This example shows how to setup service policy.

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# flowspec
RP/0/RP0/CPU0:router(config-flowspec)# address-family ipv4
RP/0/RP0/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 XR EXEC mode.

show flowspec {**afi-all** | **client** | **ipv4** | **summary** | **vrf**}

| Syntax Description | | |
|--------------------|--|------------------------------------------------------|
| afi-all | | Displays flowspec policy applied on IPv4 interfaces. |
| client | | Displays flowspec client interfaces. |
| ipv4 | | Displays flowspec policy applied on IPv4 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 XR EXEC

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | Release 5.2.4 | 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.

Examples

This example shows sample output from **show flowspec** command when **vrf**, **ipv4** and **summary** keywords are used.

```
RP/0/RP0/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
```



BFD Commands

This module provides command line interface (CLI) commands for configuring Bidirectional Forwarding Detection (BFD) on the Cisco NCS 6000 Series Router.

- [address-family ipv4 unicast \(BFD\)](#), on page 402
- [bfd](#), on page 404
- [bfd address-family ipv4 destination](#), on page 406
- [bfd address-family ipv4 fast-detect](#), on page 407
- [bfd address-family ipv4 minimum-interval](#), on page 408
- [bfd address-family ipv4 multiplier](#), on page 411
- [bfd address-family ipv4 timers](#), on page 414
- [bundle coexistence bob-blb](#), on page 416
- [bfd fast-detect](#), on page 418
- [bfd minimum-interval](#), on page 421
- [bfd multipath include location](#), on page 425
- [bfd multiplier](#), on page 427
- [clear bfd counters](#), on page 429
- [echo disable](#), on page 431
- [echo ipv4 source](#), on page 433
- [echo latency detect](#), on page 435
- [echo startup validate](#), on page 437
- [interface \(BFD\)](#), on page 439
- [ipv6 checksum](#), on page 442
- [multihop ttl-drop-threshold](#), on page 444
- [show bfd](#), on page 446
- [show bfd client](#), on page 449
- [show bfd counters](#), on page 451
- [show bfd mib session](#), on page 454
- [show bfd multipath](#), on page 457
- [show bfd session](#), on page 458
- [show bfd summary](#), on page 463

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 | | |
|-----------------------------------------|------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>address</i> | | Specifies the IPv4 unicast destination address and prefix on which to enable BFD fast-detection. |
| <i>nexthop</i> | | 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 <i>interval</i> | (Optional) | Ensures that the next hop is assigned with the same hello interval. Replace <i>interval</i> with a number that specifies the interval in milliseconds. Range is from 10 through 10000. |
| multiplier <i>multiplier</i> | (Optional) | Ensures that the next hop is assigned with the same detect multiplier. Replace <i>multiplier</i> with a number that specifies the detect multiplier. Range is from 1 through 10. |

| Command Default | |
|-----------------|-----------------------|
| | <i>interval</i> : 100 |
| | <i>multiplier</i> : 3 |

| Command Modes | |
|---------------|---------------------------------|
| | Static route configuration mode |

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | 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 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/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router static
RP/0/RP0/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 418 | Enables BFD to detect failures in the path between adjacent forwarding engines. |
| show bfd, on page 446 | Displays BFD information for a specific location. |

bfd

To enter Bidirectional Forwarding Detection (BFD) configuration mode, use the **bfd** command in XR Config 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 XR Config

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | 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.

When you issue the **bfd** command in XR Config 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/RP0/CPU0:router(config)# bfd
RP/0/RP0/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 XR Config mode
  show        Show contents of configuration
```

| Task ID | Task ID | Operations |
|---------|---------|----------------|
| | bgp | read, write |
| | ospf | read, write |
| | isis | read, write |

| Task ID | Operations |
|---------|------------|
|---------|------------|

| | |
|---------|----------------|
| mpls-te | read, write |
|---------|----------------|

Examples

The following example shows how to enter BFD configuration mode:

```
RP/0/RP0/CPU0:router # configure
RP/0/RP0/CPU0:router(config) # bfd
RP/0/RP0/CPU0:router(config-bfd) #
```

Related Commands

| Command | Description |
|----------------------------------------------|-------------------------------------------------------------------------------------------|
| echo disable, on page 431 | Disables echo mode on a router or on an individual interface or bundle. |
| interface (BFD), on page 439 | Enters BFD interface configuration mode, where you can disable echo mode on an interface. |
| show bfd, on page 446 | 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 | <i>ip-address</i> 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 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. |
|-------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

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/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# interface Bundle-Ether 1
RP/0/RP0/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 407 | Enables 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 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.

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/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# interface Bundle-Ether 1
RP/0/RP0/CPU0:router(config-if)# bfd address-family ipv4 fast-detect
```

| Related Commands | Command | Description |
|------------------|------------------------------------------------------------------|----------------------------------------------------------------------------|
| | bfd address-family ipv4 destination, on page 406 | 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 | |
| Command Default | The default is 150 ms. | |
| Command Modes | Interface configuration (config-if) | |
| Command History | Release | Modification |
| | 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.

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 $I \times M$
 - 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 $(I \times M) \times M$ —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 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:

$$\text{Asynchronous rate per bundle member} = (1000 / \text{Min-interval})$$

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).

$$\text{Echo interval} = (\text{Min-interval} \times \text{Multiplier})$$

- Calculate the overall rate supported for all members on the bundle:

$$\text{Ethernet bundle rate} = (1000 / \text{Echo interval}) \times 64$$

- Add the asynchronous base rate per bundle member to find the total rate for all bundle links:

$$\text{Total bundle rate} = \text{Ethernet bundle rate} + (\text{Base asynchronous rate} \times \text{Number of links})$$

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/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# interface Bundle-Ether 1
RP/0/RP0/CPU0:router(config-if)# bfd address-family ipv4 minimum-interval 200
```

| Related Commands | Command | Description |
|------------------|-----------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| | bfd minimum-interval, on page 421 | Specifies the minimum control packet interval for BFD sessions for the corresponding BFD configuration scope. |
| | bfd address-family ipv4 multiplier, on page 411 | 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 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.

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 $I \times M$
 - 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 $(I \times M) \times M$ —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 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 = $(1000 / \text{Min-interval})$

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 = $(\text{Min-interval} \times \text{Multiplier})$

- Calculate the overall rate supported for all members on the bundle:

Ethernet bundle rate = $(1000 / \text{Echo interval}) \times 64$

- Add the asynchronous base rate per bundle member to find the total rate for all bundle links:

Total bundle rate = Ethernet bundle rate + $(\text{Base asynchronous rate} \times \text{Number of links})$

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/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# interface Bundle-Ether 1
RP/0/RP0/CPU0:router(config-if)# bfd address-family ipv4 minimum-interval 200
RP/0/RP0/CPU0:router(config-if)# bfd address-family ipv4 multiplier 3
```


| Related Commands | Command | Description |
|------------------|-----------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------|
| | bfd address-family ipv4 minimum-interval, on page 408 | Specifies the minimum interval for asynchronous mode control packets on IPv4 BFD sessions on bundle member links. |
| | bfd minimum-interval, on page 421 | 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.

```
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. |
|--------------------|--------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| | 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. |

Command Default No timers are configured.

Command Modes Interface configuration (config-if)

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | Release 5.0.1 | 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.

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/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# interface Bundle-Ether 1
RP/0/RP0/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/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# interface Bundle-Ether 1
RP/0/RP0/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 5.2.5 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task 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/RP0/CPU0:router#configure
RP/0/RP0/CPU0:router(config)#bfd
RP/0/RP0/CPU0:router(config-bfd)#bundle coexistence bob-blb inherit
```

This example shows how to configure the coexistence mechanism between BFD over Bundle (BoB) and BFD over Logical Bundle (BLB) as being "logical":

```
RP/0/RP0/CPU0:router#configure
RP/0/RP0/CPU0:router(config)#bfd
RP/0/RP0/CPU0:router(config-bfd)#bundle coexistence bob-blb logical
```

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: .

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
 Interface configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 5.0.0 | This command was introduced. |

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 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: . In OSPF 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 |

Examples

The following example shows how to configure BFD on a BGP router:

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router bgp 65000
RP/0/RP0/CPU0:router(config-bgp)# neighbor 192.168.70.24
RP/0/RP0/CPU0:router(config-bgp-nbr)# remote-as 2
RP/0/RP0/CPU0:router(config-bgp-nbr)# bfd fast-detect
```

| Related Commands | Command | Description |
|------------------|----------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------|
| | address-family ipv4 unicast (BFD), on page 402 | Enables BFD fast-detection on a specific IPV4 unicast destination address prefix and on the forwarding next-hop address. |
| | bfd minimum-interval, on page 421 | Specifies the minimum control packet interval for BFD sessions for the corresponding BFD configuration scope. |
| | bfd multiplier, on page 427 | Sets the BFD multiplier. |

| Command | Description |
|---------------------------------------|---------------------------------------------------|
| show bfd, on page 446 | 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 | <i>milliseconds</i> 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 <i>interval</i> : 50 milliseconds IS-IS <i>interval</i> : 150 milliseconds OSPF and OSPFv3 <i>interval</i> : 150 milliseconds MPLS-TE <i>interval</i> : 15 milliseconds PIM <i>interval</i> : 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 | <table border="1"> <thead> <tr> <th>Release</th> <th>Modification</th> </tr> </thead> <tbody> <tr> <td>Release 5.0.0</td> <td>This command was introduced.</td> </tr> </tbody> </table> | Release | Modification | Release 5.0.0 | This command was introduced. |
| Release | Modification | | | | |
| Release 5.0.0 | This command was introduced. | | | | |
| Usage Guidelines | <p>To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.</p> <p>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:</p> | | | | |

- 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 7000 pps.

- If a session is running in asynchronous mode without echo, then PPS used for this session is $(1000 / \text{asynchronous interval in milliseconds})$.
- If a session is running in asynchronous mode with echo, then PPS used for this session is $(1000 / \text{echo interval in milliseconds})$.

This is calculated as: $1000 / \text{value of the } \mathbf{bfd\ minimum-interval}$ 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 |

Examples

The following example shows how to set the BFD minimum interval for a BGP routing process:

```
RP/0/RP0/CPU0:router(config)# router bgp 6500
RP/0/RP0/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 402 | Enables BFD fast-detection on a specific IPV4 unicast destination address prefix and on the forwarding next-hop address. |
| | bfd fast-detect, on page 418 | Enables BFD to detect failures in the path between adjacent forwarding engines. |
| | bfd multiplier, on page 427 | Sets the BFD multiplier. |
| | show bfd, on page 446 | Displays BFD information for a specific location. |

bfd multipath include location

To include specific linecards to host BFD multiple path sessions, use the **bfd multipath include location** command in the XR Config 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 <i>node-id</i> Configures BFD multipath on the specified location. The <i>node-id</i> variable is mentioned in the <i>rack/slot/module</i> notation. | |
| Command Default | No default behavior or values | |
| Command Modes | XR Config | |
| Command History | Release | Modification |
| | 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 |
| | 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/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# bfd multipath include location 0/5/CPU0
```

| | | |
|-------------------------|----------------------------------|--------------------------------|
| Related Commands | Command | Description |
| | bfd, on page 404 | Enters BFD configuration mode. |

| Command | Description |
|-------------------------------------------------|--------------------------------------------------------|
| show bfd multipath, on page 457 | 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
 Interface configuration
 Tunnel configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| 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.

In OSPF 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 .

| Task ID | Task ID | Operations |
|---------|---------|----------------|
| | bgp | read, write |
| | isis | read, write |
| | mpls-te | read, write |
| | ospf | read, write |

Examples

The following example shows how to set the BFD multiplier in a BGP routing process:

```
RP/0/RP0/CPU0:router(config)# router bgp 65000
RP/0/RP0/CPU0:router(config-bgp)# bfd multiplier 2
```

Related Commands

| Command | Description |
|----------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------|
| address-family ipv4 unicast (BFD), on page 402 | Enables BFD fast-detection on a specific IPV4 unicast destination address prefix and on the forwarding next-hop address. |
| bfd fast-detect, on page 418 | Enables BFD to detect failures in the path between adjacent forwarding engines. |
| bfd minimum-interval, on page 421 | Specifies the minimum control packet interval for BFD sessions for the corresponding BFD configuration scope. |
| show bfd, on page 446 | Displays BFD information for a specific location. |

clear bfd counters

To clear Bidirectional Forwarding Detection (BFD) counters, use the **clear bfd counters** command in XR 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. |
| <i>type</i> | Specifies the interface type. For more information, use the question mark (?) online help function. |
| <i>interface-path-id</i> | 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 <i>node-id</i> | Clears BFD counters from the specified location. The <i>node-id</i> argument is entered in the <i>rack/slot/module</i> 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 XR EXEC

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | 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.

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/RP0/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/RP0/CPU0:router# clear bfd counters ipv4 timing location 0/5/cpu0
```

Related Commands

| Command | Description |
|---------------------------------------------------|---------------------------------------------------------------------------------------------------------------|
| bfd fast-detect, on page 418 | Enables BFD to detect failures in the path between adjacent forwarding engines. |
| bfd minimum-interval, on page 421 | Specifies the minimum control packet interval for BFD sessions for the corresponding BFD configuration scope. |
| bfd multiplier, on page 427 | Sets the BFD multiplier. |
| show bfd, on page 446 | Displays BFD information for a specific location. |

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 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 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.

| 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/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# bfd
RP/0/RP0/CPU0:router(config-bfd)# echo disable
```

The following example shows how to disable echo mode on an individual interface:

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# bfd
RP/0/RP0/CPU0:router(config-bfd)# interface gigabitethernet 0/1/0/0
RP/0/RP0/CPU0:router(config-bfd-if)# echo disable
```

Related Commands

| Command | Description |
|-------------------------------------------------|---------------------------------------------------------------|
| bfd, on page 404 | Enters BFD configuration mode. |
| interface (BFD), on page 439 | Enters BFD interface configuration mode. |
| ipv4 verify unicast source reachable-via | Enables and disables IPv4 uRPF checking on an IPv4 interface. |
| show bfd, on page 446 | 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 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 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/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# bfd
RP/0/RP0/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/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# bfd
RP/0/RP0/CPU0:router(config-bfd)# interface gigabitethernet 0/1/0/0
RP/0/RP0/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/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# bfd
RP/0/RP0/CPU0:router(config-bfd)# interface pos 0/1/0/0
RP/0/RP0/CPU0:router(config-bfd-if)# echo ipv4 source 10.10.10.1
```

Related Commands

| Command | Description |
|-------------------------------------------|-------------------------------------------------------------------------|
| bfd, on page 404 | Enters BFD configuration mode. |
| echo disable, on page 431 | 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 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.



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 ($I \times M \times \%$), and the roundtrip delay is computed for the echo packet. If this delay is greater than ($I \times M \times \%$), 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/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router (config)# bfd
RP/0/RP0/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/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router (config)# bfd
RP/0/RP0/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/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router (config)# bfd
RP/0/RP0/CPU0:router (config-bfd)# echo latency detect percentage 100 count 3
```

Related Commands

| Command | Description |
|----------------------------------------------------|---------------------------------------------------------------------------------------------------------------|
| bfd, on page 404 | Enters BFD configuration mode. |
| bfd minimum-interval, on page 421 | Specifies the minimum control packet interval for BFD sessions for the corresponding BFD configuration scope. |
| bfd multiplier, on page 427 | Sets the BFD multiplier. |
| echo startup validate, on page 437 | 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 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. |
|-------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|



| | |
|-------------|------------------------------------------------------------------------------------------------------------------------|
| 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/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# bfd
RP/0/RP0/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/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# bfd
RP/0/RP0/CPU0:router(config-bfd)# echo startup validate force
```

Related Commands

| Command | Description |
|--------------------------------------------------|-------------------------------------------------|
| bfd, on page 404 | Enters BFD configuration mode. |
| echo latency detect, on page 435 | Enables latency detection for BFD echo packets. |

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 | <i>type</i> | Interface type. For more information, use the question mark (?) online help function. |
| | <i>interface-path-id</i> | 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 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.

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/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router (config)# bfd
RP/0/RP0/CPU0:router (config-bfd)# interface gigabitethernet 0/1/0/0
RP/0/RP0/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/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router (config)# bfd
RP/0/RP0/CPU0:router (config-bfd)# interface pos 0/1/0/0
RP/0/RP0/CPU0:router (config-bfd-if)#
```

Related Commands

| Command | Description |
|-----------------------------------------------------|------------------------------------------------------------------------|
| bfd, on page 404 | Enters BFD configuration mode. |
| echo disable, on page 431 | 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. |
| ipv6 verify unicast source reachable-via any | Enables and disables loose IPv6 uRPF checking on an IPv6 interface. |

| Command | Description |
|---------------------------------------|---------------------------|
| show bfd, on page 446 | 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 5.0.0 | This command was introduced. |
| Usage Guidelines | <p>To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.</p> <p>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.</p> | |
| 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/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# bfd
RP/0/RP0/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/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# bfd
RP/0/RP0/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/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# bfd
RP/0/RP0/CPU0:router(config-bfd)# interface gigabitethernet 0/1/0/0
RP/0/RP0/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/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# bfd
RP/0/RP0/CPU0:router(config-bfd)# interface gigabitethernet 0/1/0/0
RP/0/RP0/CPU0:router(config-bfd-if)# ipv6 checksum disable
```

Related Commands

| Command | Description |
|----------------------------------------------|-------------------------------------------------------------------------------------------|
| bfd, on page 404 | Enters BFD configuration mode. |
| interface (BFD), on page 439 | Enters BFD interface configuration mode, where you can disable echo mode on an interface. |
| show bfd, on page 446 | 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 | <i>value</i> 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 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. |
|-------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

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/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# bfd multihop ttl-drop-threshold 2
```

Related Commands

| Command | Description |
|-------------------------------------------------|--------------------------------------------------------|
| show bfd counters, on page 451 | Displays BFD counter information. |
| show bfd multipath, on page 457 | 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 XR EXEC mode.

```
show bfd [{ipv4 | [{singlehop | multihop }]} | ipv6 [{singlehop | }]] all[label] interface[ {destination | } ] [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 <i>node-id</i> argument is entered in the <i>rack/slot/module</i> 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

XR EXEC

Command History

| Release | Modification |
|---------------|------------------------------|
| 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 | Operations |
|---------|------------|
| bgp | read |
| ospf | read |
| isis | read |

| Task ID | Operations |
|---------|------------|
| mpls-te | read |

Examples

The following example shows the output from the **show bfd** command:

```
RP/0/RP0/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/RP0/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/RP0/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/RP0/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/RP0/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: 1001:1:1:1:1:1:2
Src: 1001:1:1:1:1:1:1
  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 418 | Enables BFD to detect failures in the path between adjacent forwarding engines. |
| bfd minimum-interval, on page 421 | Specifies the minimum control packet interval for BFD sessions for the corresponding BFD configuration scope. |
| bfd multiplier, on page 427 | Sets the BFD multiplier. |

show bfd client

To display Bidirectional Forwarding Detection (BFD) client information, use the **show bfd client** command in XR 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 | XR EXEC | | | | |
| Command History | <table border="1"> <thead> <tr> <th>Release</th> <th>Modification</th> </tr> </thead> <tbody> <tr> <td>Release 5.0.0</td> <td>This command was introduced.</td> </tr> </tbody> </table> | Release | Modification | Release 5.0.0 | This command was introduced. |
| Release | Modification | | | | |
| 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 | 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/RP0/CPU0:router# show bfd client

Name           Node           Num sessions
-----
bgp             0//CPU0 0
isis           0//CPU0 0
isis           0//CPU0 0
```

Table 39: 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 418 | Enables BFD to detect failures in the path between adjacent forwarding engines. |
| bfd minimum-interval, on page 421 | Specifies the minimum control packet interval for BFD sessions for the corresponding BFD configuration scope. |
| bfd multiplier, on page 427 | Sets the BFD multiplier. |
| show bfd, on page 446 | 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 XR EXEC mode.

```
show bfd counters [{ipv4 | [{singlehop | multihop}]] | ipv6 [{singlehop | }]] 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. |
| <i>type</i> | Interface type. For more information, use the question mark (?) online help function. |
| <i>interface-path-id</i> | 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 <i>node-id</i> argument is entered in the <i>rack/slot/module</i> 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

XR EXEC

Command History

| Release | Modification |
|---------------|------------------------------|
| 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.

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/RP0/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      Echo:      Recv      Xmit
  Async:             520       515      9400      9400
```

```
IPv6:
```

```
-----
```

```
POS 0/1/0/0          Recv      Xmit      Echo:      Recv      Xmit
  Async:             237       237       0         0
```

Table 40: 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 418 | Enables BFD to detect failures in the path between adjacent forwarding engines. |

| Command | Description |
|---------------------------------------------------|---------------------------------------------------------------------------------------------------------------|
| bfd minimum-interval, on page 421 | Specifies the minimum control packet interval for BFD sessions for the corresponding BFD configuration scope. |
| bfd multiplier, on page 427 | Sets the BFD multiplier. |

show bfd mib session

To display IPv4 and IPv6 Bidirectional Forwarding Detection (BFD) MIB session information, use the **show bfd mib session** command in XR EXEC mode.

show bfd mib session [**location** *node-id*]

| | | |
|---------------------------|-----------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax Description | location <i>node-id</i> | (Optional) Displays all IPv4 and IPv6 BFD MIB session information stored on the specified node. The <i>node-id</i> argument is entered in the <i>rack/slot/module</i> notation. |
|---------------------------|-----------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

| | |
|------------------------|---------------------------------------------------------------------------------------------------------------------------------------------|
| Command Default | When <i>node-id</i> is not specified, information for all IPv4 and IPv6 BFD MIB sessions, stored on the route processor node, is displayed. |
|------------------------|---------------------------------------------------------------------------------------------------------------------------------------------|

| | |
|----------------------|---------|
| Command Modes | XR EXEC |
|----------------------|---------|

| | | |
|------------------------|----------------|------------------------------|
| Command History | Release | Modification |
| | 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.

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/RP0/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/RP0/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 41: 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 <i>seconds.nanoseconds</i> , 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 <i>seconds.nanoseconds</i> , 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 458 | |

show bfd multipath

To display information concerning only BFD multipath sessions, use the **show bfd multipath** command in the XR 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 <i>node-id</i> | Displays BFD counters from the specified location. The <i>node-id</i> argument is entered in the <i>rack/slot/module</i> notation. |

Command Default No default behavior or values

Command Modes XR EXEC

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | 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 |
|---------|---------|-----------|
| | bgp | read |
| | ospf | read |
| | isis | read |
| | mpls-te | read |

show bfd session

To display Bidirectional Forwarding Detection (BFD) session information, use the **show bfd session** command in XR 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. |
| <i>type</i> | Interface type. For more information, use the question mark (?) online help function. |
| <i>interface-path-id</i> | 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 <i>node-id</i> argument is entered in the <i>rack/slot/module</i> 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 XR EXEC

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | 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.

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 session** command with the **detail** keyword and IPv4 as the default:

```
RP/0/RP0/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/RP0/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 42: show bfd session detail command Field Descriptions

| Field | Description |
|--------------------|----------------------------------------------------------------------------------------------------------------------------|
| I/f | Interface type. |
| Location | Location of the node that hosts the local endpoint of the connection, in the <i>rack/slot/module</i> 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. |

| Field | Description |
|------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Received parameters | <p>Provides information on the last transmitted control packet for the session:</p> <ul style="list-style-type: none"> • 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 | <p>Provides information on the last transmitted control packet for the session:</p> <ul style="list-style-type: none"> • 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. |
| Timer Values | <p>Provides information on the timer values used by the local and remote ends, as follows:</p> <ul style="list-style-type: none"> • 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. |

| Field | Description |
|---------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Local Stats | <p>Displays the local transmit and receive statistics,</p> <ul style="list-style-type: none"> Intervals between async packets—provides measurements on intervals between control packets (tx and rx): <ul style="list-style-type: none"> 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: <ul style="list-style-type: none"> 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. |
| Session owner information | <p>Provides the following information about the session owner.</p> <ul style="list-style-type: none"> 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 418 | Enables BFD to detect failures in the path between adjacent forwarding engines. |
| bfd multiplier, on page 427 | Sets the BFD multiplier. |
| show bfd mib session, on page 454 | 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 XR EXEC mode.

show bfd summary [{private}]location*node-id*

| Syntax Description | | |
|--------------------------------|--|------------------------------------------------------------------------------------------------------------------------------------|
| private | | Displays the private information. |
| location <i>node-id</i> | | Displays BFD counters from the specified location. The <i>node-id</i> argument is entered in the <i>rack/slot/module</i> notation. |

Command Default No default behavior or values

Command Modes XR EXEC

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | 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 |
|---------|---------|-----------|
| | 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/RP0/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:

show bfd summary

```
RP/0/RP0/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* in *Routing Configuration Guide for Cisco NCS 6000 Series Routers*.

- [address-family \(EIGRP\)](#), on page 467
- [authentication keychain](#), on page 468
- [auto-summary \(EIGRP\)](#), on page 469
- [bandwidth-percent \(EIGRP\)](#), on page 470
- [clear eigrp neighbors](#), on page 471
- [clear eigrp topology](#), on page 472
- [default-information](#), on page 473
- [default-metric \(EIGRP\)](#), on page 474
- [distance \(EIGRP\)](#), on page 475
- [hello-interval \(EIGRP\)](#), on page 477
- [hold-time \(EIGRP\)](#), on page 478
- [interface \(EIGRP\)](#), on page 480
- [log-neighbor-changes](#), on page 482
- [log-neighbor-warnings](#), on page 483
- [maximum-paths \(EIGRP\)](#), on page 484
- [metric \(EIGRP\)](#), on page 485
- [metric maximum-hops](#), on page 487
- [metric weights](#), on page 488
- [neighbor \(EIGRP\)](#), on page 490
- [next-hop-self disable](#), on page 492
- [nsf disable \(EIGRP\)](#), on page 493
- [passive-interface \(EIGRP\)](#), on page 494
- [redistribute \(EIGRP\)](#), on page 495
- [route-policy \(EIGRP\)](#), on page 497
- [router eigrp](#), on page 498
- [router-id \(EIGRP\)](#), on page 499
- [show protocols \(EIGRP\)](#), on page 500
- [split-horizon disable \(EIGRP\)](#), on page 503
- [stub \(EIGRP\)](#), on page 504

- [summary-address \(EIGRP\)](#), on page 506
- [timers active-time](#), on page 508
- [timers nsf route-hold \(EIGRP\)](#), on page 509
- [variance](#), on page 510

address-family (EIGRP)

To enable an IPv4 or IPv6 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 | <p>ipv4 Selects IPv4 address family.</p> <p>ipv6 Selects IPv6 address family.</p> | | | | |
|---------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------|--------------|---------------|------------------------------|
| Command Default | No default behavior or values | | | | |
| Command Modes | Router configuration | | | | |
| Command History | <table border="1"> <thead> <tr> <th>Release</th> <th>Modification</th> </tr> </thead> <tbody> <tr> <td>Release 5.0.0</td> <td>This command was introduced.</td> </tr> </tbody> </table> | Release | Modification | Release 5.0.0 | This command was introduced. |
| Release | Modification | | | | |
| Release 5.0.0 | This command was introduced. | | | | |
| Usage Guidelines | <p>To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.</p> <p>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.</p> | | | | |
| Task ID | <table border="1"> <thead> <tr> <th>Task ID</th> <th>Operations</th> </tr> </thead> <tbody> <tr> <td>eigrp</td> <td>read, write</td> </tr> </tbody> </table> | Task ID | Operations | eigrp | read, write |
| Task ID | Operations | | | | |
| eigrp | read, write | | | | |

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 | <i>key-chain-name</i> Name of the authentication keychain |
|---------------------------|-----------------------------------------------------------|

| | |
|------------------------|-----------------------------|
| Command Default | Authentication is disabled. |
|------------------------|-----------------------------|

| | |
|----------------------|--------------------------------------------------------------------------------------------|
| Command Modes | IPv4 address family interface configuration IPv6 address family interface configuration |
|----------------------|--------------------------------------------------------------------------------------------|

| | |
|------------------------|--------------------------------------------|
| Command History | Release Modification |
| | Release 3.9.0 No modification. |
| | 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. |
|-------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

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.

| | |
|----------------|----------------------------------|
| Task ID | Task ID Operations |
| | eigrp read, write |

| | |
|-----------------|-----------------------------------------------------------------------------|
| Examples | The following example shows how to enable an EIGRP authentication keychain: |
|-----------------|-----------------------------------------------------------------------------|

```
RP/0/RP0/CPU0:router# configure eigrp 1
RP/0/RP0/CPU0:router (config-eigrp) # address-family ipv4
RP/0/RP0/CPU0:router (config-eigrp-af) # interface GigabitEthernet 0/2/0/0
RP/0/RP0/CPU0:router (config-eigrp-if) # authentication key chain key1
```

| | | |
|-------------------------|-------------------------------------------|--------------------------------------------------------------------------------------------------------------------------|
| Related Commands | Command | Description |
| | router eigrp, on page 498 | 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 | |
| Command History | Release | Modification |
| | Release 5.0.0 | This command was introduced. |
| Usage Guidelines | <p>To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.</p> <p>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.</p> <p>EIGRP summary routes are given an administrative distance value of 5. You cannot configure this value.</p> | |
| Task ID | Task ID | Operations |
| | eigrp | read, write |
| Related Commands | Command | Description |
| | summary-address (EIGRP), on page 506 | Configures a summary aggregate address for a specified interface. |

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 | <i>percent</i> : 50 |
|------------------------|---------------------|

| | |
|----------------------|-------------------------|
| Command Modes | Interface configuration |
|----------------------|-------------------------|

| Command History | Release | Modification |
|------------------------|----------------|------------------------------|
| | 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. |
|-------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

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/RP0/CPU0:router(config)# router eigrp 1
RP/0/RP0/CPU0:router(config-eigrp)# address-family ipv4
RP/0/RP0/CPU0:router(config-eigrp-af)# router-id 10.1.1.1
RP/0/RP0/CPU0:router(config-eigrp-af)# interface GigabitEthernet 0/1/0/0
RP/0/RP0/CPU0:router(config-eigrp-af-if)# bandwidth-percent 75
```

| Related Commands | Command | Description |
|-------------------------|-----------------------|------------------------------------------|
| | bandwidth (interface) | Sets a bandwidth value for an interface. |

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 XR EXEC mode.

```
clear eigrp [as-number] [{ipv4 | ipv6}] neighbors [{ip-addresstype interface-path-id}] [soft]
```

| Syntax Description | |
|--------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>as-number</i> | (Optional) Autonomous system number. Range is from 1 to 65535. |
| ipv4 | (Optional) Specifies the IPv4 address family. |
| ipv6 | (Optional) Specifies the IPv6 address family. |
| <i>ip-address</i> | (Optional) Address of the neighbor. |
| <i>type</i> | Interface type. For more information, use the question mark (?) online help function. |
| <i>interface-path-id</i> | 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 Modes | |
|---------------|---------|
| | XR EXEC |

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | 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 | Operations |
|---------|---------|----------------|
| | eigrp | read, write |

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 XR EXEC mode.

```
clear eigrp [as-number] [{ipv4 | ipv6}] topology [{prefix mask prefixl/ength}]
```

Syntax Description

| | |
|------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| as-number | (Optional) Autonomous system number. Range is from 1 to 65535. |
| ipv4 | (Optional) Specifies the IPv4 address family. |
| ipv6 | (Optional) Specifies the IPv6 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 <i>length</i> is used, the slash is required. |

Command Default

No EIGRP topology entries are cleared.

Command Modes

XR EXEC

Command History

| Release | Modification |
|---------------|------------------------------|
| 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 | Operations |
|---------|----------------|
| eigrp | read, write |

Examples

The following example shows how to clear EIGRP topology entries for a specific route:

```
RP/0/RP0/CPU0:router# clear eigrp topology 10.1.0.0/8
```

default-information

To control the candidate default routing information for an Enhanced Interior Gateway Routing Protocol (EIGRP), use the **default-information** 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 | |
| Command History | Release | Modification |
| | 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 | Operations |
| | eigrp | read, write |

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 | |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>bandwidth</i> | Minimum bandwidth of the route in kilobits per second. Range is 1 to 4294967295. |
| <i>delay</i> | Route delay in ten microsecond units. Range is 1 to 4294967295. |
| <i>reliability</i> | 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. |
| <i>loading</i> | Effective bandwidth of the route expressed as a number from 1 to 255 (255 is 100-percent loading). |
| <i>mtu</i> | 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
 IPv6 address family configuration

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | 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.

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 |

| 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 | <p><i>internal-distance</i> 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.</p> <p><i>external-distance</i> 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.</p> | |
| Command Default | <p><i>internal-distance</i> : 90</p> <p><i>external-distance</i> : 170</p> | |
| Command Modes | <p>IPv4 address family configuration</p> <p>IPv6 address family configuration</p> | |
| Command History | Release | Modification |
| | Release 5.0.0 | This command was introduced. |
| Usage Guidelines | <p>To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.</p> <p>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.</p> <p>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.</p> <p>To display the default administrative distance for a specified routing process, use the show protocols EXEC command.</p> | |
| Task ID | Task ID | Operations |
| | eigrp | read, write |

Related Commands

| Command | Description |
|-----------------------------------------------------|----------------------------------------------------------------------------------------------------------|
| show protocols (EIGRP), on page 500 | 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 | <i>seconds</i> 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 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 | Operations |
| | eigrp | read, write |

Examples

The following example shows how to set the hello interval to 10 seconds for the interface:

```
RP/0/RP0/CPU0:router(config)# router eigrp 1
RP/0/RP0/CPU0:router(config-eigrp)# address-family ipv4
RP/0/RP0/CPU0:router(config-eigrp)# router-id 10.1.1.1
RP/0/RP0/CPU0:router(config-eigrp-af)# interface GigabitEthernet 0/1/0/0
RP/0/RP0/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 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.

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/RP0/CPU0:router(config)# router eigrp 1
RP/0/RP0/CPU0:router(config-eigrp)# address-family ipv4
RP/0/RP0/CPU0:router(config-eigrp)# router-id 10.1.1.1
RP/0/RP0/CPU0:router(config-eigrp-af)# interface GigabitEthernet 0/1/0/0
RP/0/RP0/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 477 | 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

IPv6 address family configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| 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.

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/RP0/CPU0:router(config)# router eigrp 1  
RP/0/RP0/CPU0:router(config-eigrp)# address-family ipv4  
RP/0/RP0/CPU0:router(config-eigrp)# router-id 10.1.1.1  
RP/0/RP0/CPU0:router(config-eigrp-af)# interface GigabitEthernet 0/1/0/0  
RP/0/RP0/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
 IPv6 address family configuration

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | 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.

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/RP0/CPU0:router(config)# router eigrp 1
RP/0/RP0/CPU0:router(config-eigrp)# address-family ipv4
RP/0/RP0/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
 IPv6 address family configuration

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | 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.

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 |

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

IPv6 address family configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| 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.

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 |

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 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.

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/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router eigrp 100
RP/0/RP0/CPU0:router(config-eigrp)# address-family ipv4
RP/0/RP0/CPU0:router(config-eigrp-af)# interface GigabitEthernet 0/1/0/1
```

metric (EIGRP)

```
RP/0/RP0/CPU0:router(config-eigrp-af-if)# metric bandwidth 100 delay 7 reliability 250 load 100
```

Related Commands

| Command | Description |
|-----------------------------------------------------|-------------------------------------------------------------------------|
| default-metric (EIGRP), on page 474 | 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 | <i>hops-number</i> : 100 |
|------------------------|--------------------------|

| | |
|----------------------|------------------------------------------------------------------------|
| Command Modes | IPv4 address family configuration IPv6 address family configuration |
|----------------------|------------------------------------------------------------------------|

| Command History | Release | Modification |
|------------------------|----------------|------------------------------|
| | 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. |
|-------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

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/RP0/CPU0:router(config)# router eigrp 1
RP/0/RP0/CPU0:router(config-eigrp) address-family ipv4
RP/0/RP0/CPU0:router(config-eigrp-af)# metric maximum-hops 200
```

| Related Commands | Command | Description |
|-------------------------|---------------------------------------------|-----------------------------------------------------|
| | metric weights, on page 488 | Allows the tuning of the EIGRP metric calculations. |

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 | | |
|--------------------|-----------------------|-----------------------------------------------------------------------------------------------------|
| | <i>tos</i> | Type of service (ToS) which must always be 0. |
| | <i>k1 k2 k3 k4 k5</i> | Constants that convert an EIGRP metric vector into a scalar quantity. The range is 0 to 4294967295. |

| Command Default | |
|-----------------|----------------|
| | <i>tos</i> : 0 |
| | <i>k1</i> : 1 |
| | <i>k2</i> : 0 |
| | <i>k3</i> : 1 |
| | <i>k4</i> : 0 |
| | <i>k5</i> : 0 |

| Command Modes | |
|---------------|-----------------------------------|
| | IPv4 address family configuration |
| | IPv6 address family configuration |

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | 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.

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:

$$\text{metric} = [k1 * \text{bandwidth} + (k2 * \text{bandwidth}) / (256 - \text{load}) + k3 * \text{delay}]$$

If *k5* does not equal zero, an additional operation is performed:

$$\text{metric} = \text{metric} * [k5 / (\text{reliability} + k4)]$$

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 43: 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/RP0/CPU0:router(config)# router eigrp 1
RP/0/RP0/CPU0:router(config-eigrp) address-family ipv4
RP/0/RP0/CPU0:router(config-eigrp-af)# metric weights 0 2 0 2 0 0
```

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 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.

Use the **neighbor** command to permit the point-to-point (nonbroadcast) exchange of routing information.

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/RP0/CPU0:router(config)# router eigrp 100
RP/0/RP0/CPU0:router(config-eigrp) address-family ipv4
RP/0/RP0/CPU0:router(config-eigrp-af)# interface GigabitEthernet 0/0/0/3
RP/0/RP0/CPU0:router(config-eigrp-af-if)# neighbor 172.20.1.2 remote 10
```

Related Commands

| Command | Description |
|--------------------------------------------------------|-----------------------------------------------------------------------|
| passive-interface (EIGRP), on page 494 | Disables sending and receiving "hello" messages on (EIGRP) interface. |

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 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.

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/RP0/CPU0:router(config)# router eigrp 1
RP/0/RP0/CPU0:router(config-eigrp) address-family ipv4
RP/0/RP0/CPU0:router(config-eigrp-af)# interface GigabitEthernet 0/1/0/0
RP/0/RP0/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 nsf disable** 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 Configuration
 IPv4 address family configuration
 IPv6 address family configuration

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | 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 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 |

| Related Commands | Command | Description |
|------------------|-------------------------------------------|---------------------------------------------------------------------------------------------------------------------------|
| | router eigrp, on page 498 | 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 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.

Use the **passive-interface** command to disable sending of "hello" messages. 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/RP0/CPU0:router(config)# router eigrp 1
RP/0/RP0/CPU0:router(config-eigrp) address-family ipv6
RP/0/RP0/CPU0:router(config-eigrp-af) # interface GigabitEthernet 0/6/5/0
RP/0/RP0/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 | ospfv3 | 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. |
| ospfv3 | Distributes routes from the OSPFv3 protocol. This protocol is supported in the IPv6 address family. |
| static | Redistributes IP static routes. |
| <i>as-number / instance-name</i> | 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) in 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 <i>name</i> | (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
IPv6 address family configuration

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | 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.

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.

For information about routing policies, see the *Routing Policy Commands* on module of the *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/RP0/CPU0:router(config)# router eigrp 1
RP/0/RP0/CPU0:router(config-eigrp) address-family ipv4
RP/0/RP0/CPU0:router(config-eigrp-af)# redistribute bgp 100
```

| Related Commands | Command | Description |
|------------------|-----------------------------------------------------|----------------------------|
| | default-metric (EIGRP), on page 474 | Sets metrics for an EIGRP. |

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 | |
|--------------------|------------------------------------------------|
| | <i>route-policy-name</i> 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
IPv6 address family configuration
Interface configuration

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | 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.

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/RP0/CPU0:router(config)# router eigrp 1
RP/0/RP0/CPU0:router(config-eigrp) address-family ipv4
RP/0/RP0/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 XR Config 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 | <i>instance-autonomous-system-number</i> 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 | XR Config | |
| Command History | Release | Modification |
| | Release 5.0.0 | This command was introduced. |
| Usage Guidelines | <p>To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.</p> <ul style="list-style-type: none"> • 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/RP0/CPU0:router(config)# router eigrp 109
RP/0/RP0/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 | <i>router-id</i> 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 IPv6 address family configuration | |
| Command History | Release | Modification |
| | Release 5.0.0 | This command was introduced. |
| Usage Guidelines | <p>To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.</p> <p>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.</p> | |
| 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/RP0/CPU0:router(config)# router eigrp 1
RP/0/RP0/CPU0:router(config-eigrp) address-family ipv4
RP/0/RP0/CPU0:router(config-eigrp-af)# router-id 172.20.1.1
```

show protocols (EIGRP)

To display information about the Enhanced Interior Gateway Routing Protocol (EIGRP) process configuration, use the **show protocols** command in XR EXEC mode.

```
show protocols [{ipv4 |afi-all}] [{all protocol}] [{private}]
```

| Syntax Description | |
|------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| ipv4 | (Optional) Specifies an IPv4 address family. |
| ipv6 | (Optional) Specifies an IPv6 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 . For the IPv6 address family, the options are bgp , isis , and ospfv3 . |
| default-context | (Optional) Displays default context information. 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 XR EXEC

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | 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.

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/RP0/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 44: show protocols Field Descriptions

| Field | Descriptions |
|----------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| instance | AS number of the instance. <ul style="list-style-type: none"> • 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. <ul style="list-style-type: none"> • 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. |
| Variance | Metric variance used to find feasible paths for a route. |

| Field | Descriptions |
|-----------------|---------------------------------------------------------------------------------------------------------------------------------------|
| 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. |

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 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 | Operations |
|---------|---------|----------------|
| | eigrp | read, write |

Examples

The following example shows how to disable split horizon on a GigabitEthernet link:

```
RP/0/RP0/CPU0:router(config)# router eigrp 1
RP/0/RP0/CPU0:router(config-eigrp) address-family ipv4
RP/0/RP0/CPU0:router(config-eigrp-af)# interface GigabitEthernet 0/1/0/0
RP/0/RP0/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

IPv6 address family configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| 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.

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/RP0/CPU0:router(config)# router eigrp 1
RP/0/RP0/CPU0:router(config-eigrp) address-family ipv4
RP/0/RP0/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/RP0/CPU0:router(config)# router eigrp 1
RP/0/RP0/CPU0:router(config-eigrp) address-family ipv4
RP/0/RP0/CPU0:router(config-eigrp-af)# stub receive-only
```

Related Commands

| Command | Description |
|------------------------------------------------------|-------------------------------------------------------------------------------------------------|
| redistribute (EIGRP), on page 495 | Redistributes routes from one routing domain into EIGRP. |
| summary-address (EIGRP), on page 506 | Configures a summary aggregate address for the specified EIGRP interface. |
| auto-summary (EIGRP), on page 469 | 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 | | |
|-----------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--|
| <i>ip-address</i> | The IP address argument specifies the summary IP address to apply to an interface in four-part, dotted-decimal notation. | |
| <i>/ length</i> | 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 <i>length</i> is used, the slash is required. | |
| <i>mask</i> | IP address mask. | |
| <i>admin-distance</i> | (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 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.

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/RP0/CPU0:router(config)# router eigrp 1  
RP/0/RP0/CPU0:router(config-eigrp) address-family ipv4  
RP/0/RP0/CPU0:router(config-eigrp-af)# interface GigabitEthernet 0/1/0/0  
RP/0/RP0/CPU0:router(config-eigrp-af-if)# summary-address 192.168.0.0/16 95
```

Related Commands

| Command | Description |
|---------------------------------------------------|-------------------------------------------------------------------------------------------------|
| auto-summary (EIGRP), on page 469 | 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 | <i>time-limit</i> 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 IPv6 address family configuration | |
| Command History | Release | Modification |
| | Release 5.0.0 | This command was introduced. |
| Usage Guidelines | <p>To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.</p> <p>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.</p> | |
| Task ID | Task ID | Operations |
| | eigrp | read, write |
| Examples | <p>The following example shows how to configure an indefinite routing wait time on the specified EIGRP route:</p> <pre>RP/0/RP0/CPU0:router(config)# router eigrp 1 RP/0/RP0/CPU0:router(config-eigrp) address-family ipv4 RP/0/RP0/CPU0:router(config-eigrp-af)# timers active-time disabled</pre> | |

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 | <i>seconds</i> 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. <i>seconds</i> :480 | |
| Command Modes | IPv4 address family configuration IPv6 address family configuration | |
| Command History | Release | Modification |
| | Release 5.0.0 | This command was introduced. |
| Usage Guidelines | <p>To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.</p> <p>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.</p> | |
| Task ID | Task ID | Operations |
| | eigrp | read, write |
| Examples | <p>The following example shows how to set the route hold timer value for an NSF-aware router to 2 minutes (120 seconds):</p> <pre>RP/0/RP0/CPU0:router(config)# router eigrp 1 RP/0/RP0/CPU0:router(config-eigrp) address-family ipv4 RP/0/RP0/CPU0:router(config-eigrp-af)# timers nsf route-hold 120</pre> | |

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 | <i>multiplier</i> Metric value used for load balancing. Range is from 1 to 128. |
|---------------------------|---------------------------------------------------------------------------------|

| | |
|------------------------|---------------------------------------------------|
| Command Default | <i>multiplier</i> : 1 (equal-cost load balancing) |
|------------------------|---------------------------------------------------|

| | |
|----------------------|------------------------------------------------------------------------|
| Command Modes | IPv4 address family configuration IPv6 address family configuration |
|----------------------|------------------------------------------------------------------------|

| Command History | Release | Modification |
|------------------------|----------------|------------------------------|
| | 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. |
|-------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

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/RP0/CPU0:router(config)# router eigrp 1
RP/0/RP0/CPU0:router(config-eigrp) address-family ipv4
RP/0/RP0/CPU0:router(config-eigrp-af)# variance 4
```



IS-IS Commands

This module describes the commands used to configure and monitor the Intermediate System-to-Intermediate System (IS-IS).

For detailed information about IS-IS concepts, configuration tasks, and examples, see the *Implementing IS-IS on* module in the *Routing Configuration Guide for Cisco NCS 6000 Series Routers*.

- [address-family \(IS-IS\), on page 514](#)
- [adjacency-check disable, on page 516](#)
- [adjacency stagger, on page 518](#)
- [attached-bit receive ignore, on page 520](#)
- [attached-bit send, on page 521](#)
- [circuit-type, on page 523](#)
- [clear isis process, on page 525](#)
- [clear isis route, on page 526](#)
- [clear isis statistics, on page 528](#)
- [csnp-interval, on page 530](#)
- [default-information originate \(IS-IS\), on page 532](#)
- [disable \(IS-IS\), on page 534](#)
- [distance \(IS-IS\), on page 535](#)
- [distribute-list in \(IS-IS\), on page 537](#)
- [fast-reroute per-link \(IS-IS\), on page 539](#)
- [fast-reroute per-prefix \(IS-IS\), on page 541](#)
- [fast-reroute per-link priority-limit \(IS-IS\), on page 543](#)
- [fast-reroute per-prefix load-sharing disable \(IS-IS\), on page 544](#)
- [fast-reroute per-prefix tiebreaker \(IS-IS\), on page 545](#)
- [hello-interval \(IS-IS\), on page 547](#)
- [hello-multiplier, on page 549](#)
- [hello-padding, on page 551](#)
- [hello-password, on page 553](#)
- [hello-password keychain, on page 555](#)
- [hello-password accept, on page 557](#)
- [hostname dynamic disable, on page 559](#)
- [ignore-lsp-errors, on page 560](#)
- [instance-id, on page 561](#)
- [interface \(IS-IS\), on page 562](#)

- ipfrr lfa, on page 564
- ipfrr lfa exclude interface, on page 566
- is-type, on page 568
- link-down fast-detect, on page 570
- link-group, on page 571
- log adjacency changes (IS-IS), on page 572
- log pdu drops, on page 573
- lsp fast-flood threshold, on page 574
- lsp-gen-interval, on page 575
- lsp-interval, on page 577
- lsp-mtu, on page 578
- lsp-password, on page 580
- lsp-password accept, on page 582
- lsp-refresh-interval, on page 584
- maximum-paths (IS-IS), on page 586
- maximum-redistributed-prefixes (IS-IS), on page 587
- max-lsp-lifetime, on page 588
- mesh-group (IS-IS), on page 589
- metric (IS-IS), on page 591
- metric-style narrow, on page 593
- metric-style transition, on page 595
- metric-style wide, on page 597
- microloop avoidance, on page 599
- microloop avoidance segment-routing, on page 601
- min-lsp-arrivaltime, on page 603
- mpls ldp auto-config , on page 605
- mpls ldp sync (IS-IS), on page 606
- mpls traffic-eng (IS-IS), on page 607
- mpls traffic-eng path-selection ignore overload, on page 609
- mpls traffic-eng router-id (IS-IS), on page 611
- mpls traffic-eng srlg (IS-IS), on page 613
- net, on page 615
- nsf (IS-IS), on page 617
- nsf interface-expires, on page 619
- nsf interface-timer, on page 621
- nsf lifetime (IS-IS), on page 623
- passive (IS-IS), on page 624
- point-to-point, on page 625
- priority (IS-IS), on page 626
- propagate level, on page 628
- redistribute (IS-IS), on page 630
- retransmit-interval (IS-IS), on page 634
- retransmit-throttle-interval, on page 636
- router isis, on page 638
- route source first-hop, on page 639
- set-overload-bit, on page 640

- [set-attached-bit](#), on page 642
- [show isis](#), on page 644
- [show isis adjacency](#), on page 646
- [show isis adjacency-log](#), on page 648
- [show isis checkpoint adjacency](#), on page 650
- [show isis checkpoint interface](#), on page 652
- [show isis checkpoint lsp](#), on page 654
- [show isis database](#), on page 656
- [show isis database-log](#), on page 658
- [show isis fast-reroute](#), on page 660
- [show isis hostname](#), on page 663
- [show isis interface](#), on page 665
- [show isis lsp-log](#), on page 669
- [show isis mesh-group](#), on page 671
- [show isis mpls traffic-eng adjacency-log](#), on page 673
- [show isis mpls traffic-eng advertisements](#), on page 675
- [show isis mpls traffic-eng tunnel](#), on page 677
- [show isis neighbors](#), on page 679
- [show isis protocol](#), on page 682
- [show isis route](#), on page 684
- [show isis spf-log](#), on page 687
- [show isis statistics](#), on page 694
- [show isis topology](#), on page 698
- [show protocols \(IS-IS\)](#), on page 701
- [shutdown \(IS-IS\)](#), on page 704
- [single-topology](#), on page 705
- [snmp-server traps isis](#), on page 706
- [spf-interval](#), on page 707
- [spf prefix-priority \(IS-IS\)](#), on page 709
- [summary-prefix \(IS-IS\)](#), on page 711
- [suppressed](#), on page 713
- [tag \(IS-IS\)](#), on page 714
- [trace \(IS-IS\)](#), on page 715

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}
no address-family {ipv4 | ipv6} {unicast}
```

Syntax Description

| | |
|----------------|-------------------------------------|
| ipv4 | Specifies IPv4 address prefixes. |
| ipv6 | Specifies IPv6 address prefixes. |
| unicast | Specifies unicast 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 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.

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 IPv4 or 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/RP0/CPU0:router(config)# router isis isp
RP/0/RP0/CPU0:router(config-isis)# interface gigabitEthernet 0/1/0/0
RP/0/RP0/CPU0:router(config-isis-if)# address-family ipv4 unicast
```

```
RP/0/RP0/CPU0:router(config-isis-if-af)#
```

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 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.

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/RP0/CPU0:router(config)# router isis isp
RP/0/RP0/CPU0:router(config-isis)# address-family ipv6 |ipv4
RP/0/RP0/CPU0:router(config-isis-af)# adjacency-check disable
```


adjacency stagger

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. |
| <i>initial-num-nbr</i> | 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. |
| <i>max-num-nbr</i> | 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 45: 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 46: 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
```

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 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 | Operations |
|---------|---------|----------------|
| | isis | read, write |

Examples The following example shows how to configure to ignore the attached bit in a received LSP:

```
RP/0/RP0/CPU0:router(config)# router isis ispl
RP/0/RP0/CPU0:router(config-isis)# address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-isis-af)# attached-bit receive ignore
```

| Related Commands | Command | Description |
|------------------|------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------|
| | attached-bit send, on page 521 | 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 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.

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 630](#) 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/RP0/CPU0:router(config)# router isis ispl  
RP/0/RP0/CPU0:router(config-isis)# address-family ipv4 unicast  
RP/0/RP0/CPU0:router(config-isis-af)# attached-bit send always-set
```

Related Commands

| Command | Description |
|---------------------------------------------------|--------------------------------------------------------------------------------------------------------|
| redistribute (IS-IS), on page 630 | Redistribute routes from one routing protocol into Intermediate System-to-Intermediate System (IS-IS). |
| single-topology, on page 705 | Configures the link topology for IPv4 when IPv6 is configured. |

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 1 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 | <table border="1"> <thead> <tr> <th>Release</th> <th>Modification</th> </tr> </thead> <tbody> <tr> <td>Release 5.0.0</td> <td>This command was introduced.</td> </tr> </tbody> </table> | Release | Modification | Release 5.0.0 | This command was introduced. |
| Release | Modification | | | | |
| 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.

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 568](#) 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//CPU0:router(config)# router isis isp
RP/0//CPU0:router(config-isis)# is-type level-1-2
RP/0//CPU0:router(config-isis)# interface GigabitEthernet 0/2/0/0
RP/0//CPU0:router(config-isis-if)# circuit-type level-1
```

```
RP/0//CPU0:router(config-isis-if)# exit
RP/0//CPU0:router(config-isis)# interface GigabitEthernet 0/5/0/2
RP/0//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//CPU0:router(config)# router isis isp
RP/0//CPU0:router(config-isis)# is-type level-2-only
RP/0//CPU0:router(config-isis)# interface GigabitEthernet 0/2/0/0
RP/0//CPU0:router(config-isis-if)# circuit-type level-1-2
```

Related Commands

| Command | Description |
|--------------------------------------|----------------------------------------------------------------------------|
| is-type, on page 568 | Configures the routing level for an instance of the IS-IS routing process. |
| net, on page 615 | 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 XR EXEC 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

XR EXEC

Command History

Release

Release 5.0.0

Modification

This command was introduced.

Usage Guidelines

To use this command, you must be in a user group associated with a task group that 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/RP0/CPU0:router# clear isis instance 1 process
```

Related Commands

| Command | Description |
|--------------------------------------------------|---------------------------------------------|
| show isis database, on page 656 | Displays the IS-IS link-state database. |
| show isis neighbors, on page 679 | 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 XR EXEC mode.

clear isis [**instance** *instance-id*] {**afi-all** | **ipv4** | **ipv6**} {**unicast** | **safi-all**} [**topology** *topo-name*] **route**

Syntax Description

| | |
|------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| instance <i>instance-id</i> | (Optional) Specifies IS-IS sessions for the specified IS-IS instance only. <ul style="list-style-type: none"> The <i>instance-id</i> 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. |
| safi-all | Specifies all secondary address prefixes. |
| topology <i>topo-name</i> | (Optional) Specifies topology table information and name of the topology table. |

Command Default

No default behavior or value

Command Modes

XR EXEC

Command History

| Release | Modification |
|---------------|------------------------------|
| 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.

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 |
| basic-services | read, write |

Examples

The following example shows how to clear the routes with IPv4 unicast address prefixes:

```
RP/0/RP0/CPU0:router# clear isis ipv4 unicast route
```

Related Commands

| Command | Description |
|--------------------------------------------------|---------------------------------------------|
| show isis database, on page 656 | Displays the IS-IS link-state database. |
| show isis neighbors, on page 679 | 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 XR EXEC mode.

clear isis [**instance** *instance-id*] **statistics** [*type interface-path-id*]

| | |
|---------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax Description | <p>instance <i>instance-id</i> (Optional) Clears IS-IS sessions for the specified IS-IS instance only.</p> <ul style="list-style-type: none"> The <i>instance-id</i> argument is the instance identifier (alphanumeric) defined by the router isis command. |
| | <p><i>type</i> Interface type. For more information, use the question mark (?) online help function.</p> |
| | <p><i>interface-path-id</i> Physical interface or virtual interface.</p> <p>Note Use the show interfaces command to see a list of all interfaces currently configured on the router.</p> <p>For more information about the syntax for the router, use the question mark (?) online help function.</p> |

Command Default No default behavior or values

Command Modes XR EXEC

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | 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.

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/RP0/CPU0:router# clear isis instance 23 statistics
```

Related Commands

| Command | Description |
|---------------------------------------------------|--------------------------------|
| show isis statistics, on page 694 | Displays the IS-IS statistics. |

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 | <i>seconds</i> | 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 | <i>seconds</i> : 10 seconds |
| | Both Level 1 and Level 2 are configured if no level is specified. |

| | |
|----------------------|-------------------------|
| Command Modes | Interface configuration |
|----------------------|-------------------------|

| | | |
|------------------------|----------------|------------------------------|
| Command History | Release | Modification |
| | 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.

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/RP0/CPU0:router(config)# router isis isp
```

```
RP/0/RP0/CPU0:router(config-isis)# interface GigabitEthernet 0/0/2/0  
RP/0/RP0/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. |
| | <i>route-policy-name</i> | (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 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 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:

```
RP0/RP0/CPU0:router(config)# router isis isp
```



```
RP/0/RP0/CPU0:router(config-isis)# address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-isis-af)# default-information originate
```

Related Commands

| Command | Description |
|---------------------------------------------------|---------------------------------------------------------------------------------------------------------|
| redistribute (IS-IS), on page 630 | Redistributes routes from one routing protocol into Intermediate System-to-Intermediate System (IS-IS). |
| show isis database, on page 656 | 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 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 | 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/RP0/CPU0:router(config)# router isis isp
RP/0/RP0/CPU0:router(config-isis)# interface GigabitEthernet 0/1/0/1
RP/0/RP0/CPU0:router(config-isis-if)# address-family ipv4 unicast
RP/0/RP0/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 | |
|-------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>weight</i> | Administrative distance to be assigned to IS-IS routes. Range is 1 to 255. |
| <i>prefix</i> | (Optional) The <i>prefix</i> argument specifies the IP address in four-part, dotted-decimal notation. |
| <i>mask</i> | (Optional) IP address mask. |
| <i>/length</i> | (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. |
| <i>prefix-list-name</i> | (Optional) List of routes to which administrative distance applies. |

| Command Default | |
|-----------------|---------------------|
| | <i>weight</i> : 115 |

| Command Modes | |
|---------------|------------------------------|
| | Address family configuration |

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | 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.

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//CPU0:router(config)# ipv4 prefix-list target_routes
RP/0//CPU0:router(config-ipv4_pfx)# permit 2.0.0.0/8
RP/0//CPU0:router(config-ipv4_pfx)# permit 3.0.0.0/8
RP/0//CPU0:router(config-ipv4_pfx)# deny 0.0.0.0/0
RP/0//CPU0:router(config-ipv4_pfx)# exit
RP/0//CPU0:router(config)# router isis isp
RP/0//CPU0:router(config-isis)# address-family ipv4 unicast
RP/0//CPU0:router(config-isis-af)# distance 10 1.0.0.0/8 target_routes
RP/0//CPU0:router(config-isis-af)# distance 80
```

Related Commands

| Command | Description |
|------------------------------------------------------------|------------------------------------------------------------------------|
| router isis, on page 638 | Configures the IS-IS routing protocol and specifies an IS-IS instance. |
| show isis protocol, on page 682 | Displays summary information about the IS-IS instance. |
| show isis route, on page 684 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} | ifa-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. |
| ifa-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 ifa-candidate keyword). |
| <i>type</i> | | Interface type. For more information, use the question mark (?) online help function. |
| <i>interface-path-id</i> | | 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 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 per-link fast-reroute LFA computation for the IPv4 unicast topology at Level 1:

```
RP/0/RP0/CPU0:router(config)# router isis isp  
RP/0/RP0/CPU0:router(config-isis)# interface GigabitEthernet 0/3/0/0  
RP/0/RP0/CPU0:router(config-isis-if)# address-family ipv4 unicast  
RP/0/RP0/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). |
| <i>type</i> | | Interface type. For more information, use the question mark (?) online help function. |
| <i>interface-path-id</i> | | 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 <i>prefix-list-name</i> | | 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/RP0/CPU0:router(config)# router isis isp
RP/0/RP0/CPU0:router(config-isis)#
RP/0/RP0/CPU0:router(config-isis-if)# address-family ipv4 unicast
RP/0/RP0/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 computation for critical priority prefixes only. |
| | high | Enables LFA computation for for critical and 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

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | 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 | 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/RP0/CPU0:router#configure
RP/0/RP0/CPU0:router(config)#router isis isp_lfa
RP/0/RP0/CPU0:router(config-isis)#address-family ipv4
RP/0/RP0/CPU0:router(config-isis-af)#fast-reroute per-link priority-limit critical 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 | |
| Command History | Release | Modification |
| | 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 | Operations |
| | isis | read, write |

This example shows how to disable load-sharing prefixes across multiple backups for level 1 routes:

```
RP/0/RP0/CPU0:router#configure
RP/0/RP0/CPU0:router(config)#router isis isp_lfa
RP/0/RP0/CPU0:router(config-isis)#address-family ipv4
RP/0/RP0/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. |
| <i>index_number</i> | | 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

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | 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 | 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/RP0/CPU0:router#configure
RP/0/RP0/CPU0:router(config)#router isis isp_lfa
RP/0/RP0/CPU0:router(config-isis)#address-family ipv4
RP/0/RP0/CPU0:router(config-isis-af)#fast-reroute per-prefix tiebreaker downstream index
255
```


This example shows how to configure all the criteria for backup path selection:

```
RP/0/RP0/CPU0:router#configure
RP/0/RP0/CPU0:router(config)#router isis isp_lfa
RP/0/RP0/CPU0:router(config-isis)#address-family ipv4
RP/0/RP0/CPU0:router(config-isis-af)#fast-reroute per-prefix tiebreaker srlg-disjoint index
10
RP/0/RP0/CPU0:router(config-isis-af)#fast-reroute per-prefix tiebreaker primary-path index
20
RP/0/RP0/CPU0:router(config-isis-af)#fast-reroute per-prefix tiebreaker lowest-backup-metric
index 30
RP/0/RP0/CPU0:router(config-isis-af)#fast-reroute per-prefix tiebreaker lc-disjoint index
40
RP/0/RP0/CPU0:router(config-isis-af)#fast-reroute per-prefix tiebreaker node-protecting
index 50
```

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 | <p><i>seconds</i> Integer value (in seconds) for the length of time between consecutive hello packets. By default, a value three times the hello interval <i>seconds</i> is advertised as the <i>hold time</i> 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.</p> <p>level { 1 2 } (Optional) Specifies the hello interval for Level 1 and Level 2 independently. For broadcast interfaces only.</p> | | | | |
|-------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------|--------------|---------------|------------------------------|
| Command Default | <p><i>seconds</i> : 10 seconds</p> <p>Both Level 1 and Level 2 are configured if no level is specified.</p> | | | | |
| Command Modes | Interface configuration | | | | |
| Command History | <table border="1"> <thead> <tr> <th>Release</th> <th>Modification</th> </tr> </thead> <tbody> <tr> <td>Release 5.0.0</td> <td>This command was introduced.</td> </tr> </tbody> </table> | Release | Modification | Release 5.0.0 | This command was introduced. |
| Release | Modification | | | | |
| Release 5.0.0 | This command was introduced. | | | | |
| Usage Guidelines | <p>To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.</p> <p>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.</p> | | | | |
|  Note | <p>A shorter hello interval gives quicker convergence, but increases bandwidth and CPU usage. It might also add to instability in the network.</p> <p>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.</p> <p>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.</p> | | | | |



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/RP0/CPU0:router(config)# router isis isp
RP/0/RP0/CPU0:router(config-isis)# interface TenGigE 0/6/0/0
RP/0/RP0/CPU0:router(config-isis-if)# hello-interval 5 level 1
```

Related Commands

| Command | Description |
|-----------------------------------------------|--------------------------------------------------------------------------------------------------------------------------|
| hello-multiplier, on page 549 | 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 | <p><i>multiplier</i> 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.</p> <p>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.</p> <p>level { 1 2 } (Optional) Specifies the hello multiplier independently for Level 1 or Level 2 adjacencies.</p> | | | | |
|---------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------|--------------|---------------|------------------------------|
| Command Default | <p><i>multiplier</i> : 3</p> <p>Both Level 1 and Level 2 are configured if no level is specified.</p> | | | | |
| Command Modes | Interface configuration | | | | |
| Command History | <table border="1"> <thead> <tr> <th>Release</th> <th>Modification</th> </tr> </thead> <tbody> <tr> <td>Release 5.0.0</td> <td>This command was introduced.</td> </tr> </tbody> </table> | Release | Modification | Release 5.0.0 | This command was introduced. |
| Release | Modification | | | | |
| Release 5.0.0 | This command was introduced. | | | | |
| Usage Guidelines | <p>To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.</p> <p>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.</p> <p>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 547 command) correspondingly to make the hello protocol more reliable without increasing the time required to detect a link failure.</p> <p>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.</p> | | | | |

| 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/RP0/CPU0:router(config)# router isis isp
RP/0/RP0/CPU0:router(config-isis)# interface GigabitEthernet /2/0/1
RP/0/RP0/CPU0:router(config-isis-if)# hello-interval 6
RP/0/RP0/CPU0:router(config-isis-if)# hello-multiplier 10
```

Related Commands

| Command | Description |
|-----------------------------------------------------|-----------------------------------------------------------------------------|
| hello-interval (IS-IS), on page 547 | 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 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.

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/RP0/CPU0:router(config)# router isis isp
RP/0/RP0/CPU0:router(config-isis)# interface GigabitEthernet 0/2/0/1
RP/0/RP0/CPU0:router(config-isis-if)# hello-padding disable
```

Related Commands

| Command | Description |
|--------------------------------------------------|-------------------------------------------------|
| show isis interface, on page 665 | 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. |
| <i>password</i> | 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 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.

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/RP0/CPU0:router(config)# router isis isp
RP/0/RP0/CPU0:router(config-isis)# interface GigabitEthernet 0/2/0/3
RP/0/RP0/CPU0:router(config-isis-if)# hello-password hmac-md5 clear mypassword
```

Related Commands

| Command | Description |
|------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------|
| hello-password keychain, on page 555 | Configures the authentication password keychain for an Intermediate System-to-Intermediate System (IS-IS) interface. |
| hello-password accept, on page 557 | 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. |
|--------------------|--------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| | <i>keychain-name</i> | 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 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.

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/RP0/CPU0:routerRP/0/RP0/CPU0:router(config)# router isis isp
RP/0/RP0/CPU0:routerRP/0/RP0/CPU0:router(config-isis)# interface GigabitEthernet 0/1/0/0
```

```
RP/0/RP0/CPU0:routerRP/0/RP0/CPU0:router(config-isis-if)# hello-password keychain mykeychain  
level 1 send-only
```

Related Commands

| Command | Description |
|----------------------------------------------------|-------------------------------------------------------------------------------------------------------------|
| hello-password, on page 553 | Configures the authentication password for an Intermediate System-to-Intermediate System (IS-IS) interface. |
| hello-password accept, on page 557 | 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. |
| | <i>password</i> | 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 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. | |
| | 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: | |
| | <pre>RP/0/RP0/CPU0:routerRP/0/RP0/CPU0:router(config)# router isis isp RP/0/RP0/CPU0:routerRP/0/RP0/CPU0:router(config-isis)# interface GigabitEthernet 0/2/0/3 RP/0/RP0/CPU0:routerRP/0/RP0/CPU0:router(config-isis)# hello-password accept encrypted 111D1C1603</pre> | |

Related Commands

| Command | Description |
|---------------------------------------------|---------------------------------------------------------------|
| hello-password, on page 553 | 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 5.0.0 | 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/RP0/CPU0:router(config)# router isis isp
RP/0/RP0/CPU0:router(config-isis)# hostname dynamic disable
```

| Related Commands | Command | Description |
|-------------------------|-------------------------------------------------|------------------------------------------------------|
| | hostname | Specifies the name of the local router. |
| | show isis hostname, on page 663 | 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 | <table border="1"> <thead> <tr> <th>Release</th> <th>Modification</th> </tr> </thead> <tbody> <tr> <td>Release 5.0.0</td> <td>This command was introduced.</td> </tr> </tbody> </table> | Release | Modification | Release 5.0.0 | This command was introduced. |
| Release | Modification | | | | |
| 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.

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/RP0/CPU0:router(config)# router isis isp
RP/0/RP0/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 | <i>identifier</i> 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/RP0/CPU0:router(config)# router isis ring
RP/0/RP0/CPU0:router(config-isis)# instance-id 1
RP/0/RP0/CPU0:router(config-isis)# exit
RP/0/RP0/CPU0:router(config)# router isis 1
RP/0/RP0/CPU0:router(config-isis)# instance-id 6
RP/0/RP0/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 | <i>type</i> | Interface type. For more information, use the question mark (?) online help function. |
| | <i>interface-path-id</i> | 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 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.

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 mult topology configuration for IPv4 on GigabitEthernet interface 0/3/0/0:

```
RP/0/RP0/CPU0:router(config)# router isis isp
RP/0/RP0/CPU0:router(config-isis)# net 49.0000.0000.0001.00
RP/0/RP0/CPU0:router(config-isis)# interface GigabitEthernet 0/3/0/0
RP/0/RP0/CPU0:router(config-isis-if)# address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-isis-if-af)# metric-style wide level 1
!
```

```
RP/0/RP0/CPU0:router(config)# interface GigabitEthernet 0/3/0/0
RP/0/RP0/CPU0:router(config-if)# ipv4 address 2001::1/64
```

Related Commands

| Command | Description |
|------------------------------------------------------------|----------------------------------------------------------------------------|
| log adjacency changes (IS-IS), on page 572 | Configures the routing level for an instance of the IS-IS routing process. |
| net, on page 615 | Configures an IS-IS network entity title (NET) for the routing process. |
| router isis, on page 638 | 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 5.0.0 | 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/RP0/CPU0:router(config)# router isis isp
RP/0/RP0/CPU0:router(config-isis)# interface GigabitEthernet 0/3/0/0
RP/0/RP0/CPU0:router(config-isis-if)# address-family ipv4 unicast
RP/0/RP0/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 | <i>type</i> | Interface type. For more information, use the question mark (?) online help function. |
| | <i>interface-path-id</i> | 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 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.

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/RP0/CPU0:router(config)# router isis isp  
RP/0/RP0/CPU0:router(config-isis)# interface GigabitEthernet 0/1/0/0  
RP/0/RP0/CPU0:router(config-isis-if)# address-family ipv4 unicast  
RP/0/RP0/CPU0:router(config-isis-if-af)# ipfrr lfa exclude interface GigabitEthernet 0/1/0/0
```

Related Commands

| Command | Description |
|----------------------------------------|--------------------------------------------------------------------------|
| ipfrr lfa, on page 564 | Enable the IP fast reroute (IPFRR) loop-free alternate (LFA) computation |

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 | level-1-2 | level-2-only |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| | 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. | Specifies that the router perform both Level 1 and Level 2 routing. | 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 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.

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/RP0/CPU0:router(config)# router isis isp  
RP/0/RP0/CPU0:router(config-isis)# is-type level-2-only
```

Related Commands

| Command | Description |
|--------------------------------------------------|---------------------------------------------|
| circuit-type, on page 523 | Configures the type of adjacency. |
| show isis neighbors, on page 679 | 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

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 | <i>link-group-name</i> Name of a link group. | | | | |
|---------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------|------------|------|----------------|
| Command Default | No link groups are configured. | | | | |
| Command Modes | IS-IS interface configuration Address-family configuration | | | | |
| 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 | <table border="1"> <thead> <tr> <th>Task ID</th> <th>Operations</th> </tr> </thead> <tbody> <tr> <td>isis</td> <td>read, write</td> </tr> </tbody> </table> | Task ID | Operations | isis | read, write |
| Task ID | Operations | | | | |
| isis | read, write | | | | |
| Examples | <p>The following example shows how to configure GigabitEthernet interface 0/3/0/0 as a member of a link group:</p> <pre>RP/0/RP0/CPU0:router(config)# router isis isp RP/0/RP0/CPU0:router(config-isis)# interface GigabitEthernet 0/3/0/0 RP/0/RP0/CPU0:router(config-isis-if)# address-family ipv4 unicast RP/0/RP0/CPU0:router(config-isis-if-af)# link-group purple</pre> | | | | |

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 5.0.0 | 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/RP0/CPU0:router(config)# router isis isp
RP/0/RP0/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 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.

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_IIH_INPUT_Q_OVERFLOW: IIH input queue overflow: 86 total drops; 19 IIH drops,
44 LSP drops, 23 SNP drops
%ISIS-4-ERR_LSP_INPUT_Q_OVERFLOW: LSP input queue overflow: 17 total drops; 9 IIH drops,
3 LSP drops, 5 SNP drops
```

| Task ID | Task ID | Operations |
|---------|---------|----------------|
| | isis | read, write |

Examples The following example shows how to enable PDU logging:

```
RP/0/RP0/CPU0:router(config)# router isis isp
RP/0/RP0/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 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.

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/RP0/CPU0:router(config)# router isis isp
RP/0/RP0/CPU0:router(config-isis)# interface GigabitEthernet 0/3/0/0
RP/0/RP0/CPU0:router(config-isis-if)# lsp fast-flood threshold 234 level 1
```

Related Commands

| Command | Description |
|-------------------------------------------|------------------------------------------------------------------------------------|
| lsp-interval, on page 577 | Configures the amount of time between consecutive LSPs sent on an IS-IS interface. |

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 <i>initial</i> | | Specifies the initial LSP generation delay (in milliseconds). Range is 0 to 120000 milliseconds. |
| secondary-wait <i>secondary</i> | | Specifies the hold time between the first and second LSP generation (in milliseconds). Range is 1 to 120000 milliseconds. |
| maximum-wait <i>maximum</i> | | 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 <i>initial</i> : 50 milliseconds | |
| secondary-wait <i>secondary</i> : 200 milliseconds | |
| maximum-wait <i>maximum</i> : 5000 milliseconds | |

| Command Modes | |
|---------------|----------------------|
| | Router configuration |

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | 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.

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/RP0/CPU0:router(config)# router isis isp
RP/0/RP0/CPU0:router(config-isis)# lsp-gen-interval maximum-wait 15 initial-wait 5
```

Related Commands

| Command | Description |
|----------------------------------------------------------|--------------------------------------------------------------------------------------------------|
| retransmit-interval (IS-IS), on page 634 | 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 | <i>milliseconds</i> 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 | <i>milliseconds</i> : 33 milliseconds | | | | |
| Command Modes | Interface configuration | | | | |
| Command History | <table border="1"> <thead> <tr> <th>Release</th> <th>Modification</th> </tr> </thead> <tbody> <tr> <td>Release 5.0.0</td> <td>This command was introduced.</td> </tr> </tbody> </table> | Release | Modification | Release 5.0.0 | This command was introduced. |
| Release | Modification | | | | |
| 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 | <table border="1"> <thead> <tr> <th>Task ID</th> <th>Operations</th> </tr> </thead> <tbody> <tr> <td>isis</td> <td>read, write</td> </tr> </tbody> </table> | Task ID | Operations | isis | read, write |
| 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/RP0/CPU0:router(config)# router isis isp
RP/0/RP0/CPU0:router(config-isis)# interface GigabitEthernet /2/0/1
RP/0/RP0/CPU0:router(config-isis-if)# lsp-interval 100
```

| Related Commands | Command | Description |
|------------------|----------------------------------------------------------|--------------------------------------------------------------------------------------------------|
| | retransmit-interval (IS-IS), on page 634 | 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 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.

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 665](#) 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/RP0/CPU0:router(config)# router isis isp  
RP/0/RP0/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 665 | 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] | [enable-poi]
no lsp-password [{hmac-md5 | text}] [{clear | encrypted}] password | keychain keychain-name]
[level {1 | 2}] [send-only] [snp send-only] | [enable-poi]
```

| 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. |
| <i>password</i> | Authentication password you assign. |
| keychain | (Optional) Specifies a keychain. |
| <i>keychain-name</i> | 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. |
| enable-poi | The enable-poi keyword inserts the purge originator identification (POI), if you are using cryptographic authentication. If you are not using cryptographic authentication, then the POI is inserted by default. |

Command Default Both Level 1 and Level 2 are configured if no level is specified.

Command Modes Router configuration

| Command History | Release | Modification |
|-----------------|---------------|-----------------------------------------|
| | Release 5.0.0 | This command was introduced. |
| | Release 6.5.1 | The enable-poi keyword is added. |

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/RP0/CPU0:router(config)# router isis isp
RP/0/RP0/CPU0:router(config-isis)# lsp-password hmac-md5 clear password1 level 1
RP/0/RP0/CPU0:router(config-isis)# lsp-password text clear password2 level 2
```

Related Commands

| Command | Description |
|--------------------------------------------------|------------------------------------------------------------------------------------------------|
| lsp-password accept, on page 582 | 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. |
| | <i>password</i> | 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 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. | |
| | 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 accept Level 1 LSP and SNP password:

```
RP/0/RP0/CPU0:router(config)# router isis isp
RP/0/RP0/CPU0:router(config-isis)# lsp-password accept encrypted password1 level 1
```

Related Commands

| Command | Description |
|-------------------------------------------|--------------------------------------------|
| lsp-password, on page 580 | 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 | <i>seconds</i> | Refresh interval (in seconds). Range is 1 to 65535 seconds. |
| | level { 1 2 } | (Optional) Specifies routing Level 1 or Level 2 independently. |
| Command Default | <i>seconds</i> : 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 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.

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/RP0/CPU0:router(config)# router isis isp
RP/0/RP0/CPU0:router(config-isis)# lsp-refresh-interval 10800
```

Related Commands

| Command | Description |
|-----------------------------------------------|------------------------------------------------------------------|
| max-lsp-lifetime, on page 588 | 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 into 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 | <i>maximum</i> Maximum number of parallel routes that IS-IS can install in a routing table. Range is 1 to 32 |
|---------------------------|--------------------------------------------------------------------------------------------------------------|

| | |
|----------------------|------------------------------|
| Command Modes | Address family configuration |
|----------------------|------------------------------|

| | | |
|------------------------|----------------|------------------------------|
| Command History | Release | Modification |
| | 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 | Operations |
| | isis | read, write |

Examples The following example shows how to allow a maximum of 16 paths to a destination:

```
RP/0/RP0/CPU0:router(config)# router isis isp
RP/0/RP0/CPU0:router(config-isis)# address-family ipv4 unicast
RP/0/RP0/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 | <i>maximum</i> 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 | <i>maximum</i> : 10000 level : 1-2 |
|------------------------|----------------------------------------------|

| | |
|----------------------|------------------------------|
| Command Modes | Address family configuration |
|----------------------|------------------------------|

| | | |
|------------------------|----------------|------------------------------|
| Command History | Release | Modification |
| | 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.

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//CPU0:router(config)# router isis isp
RP/0//CPU0:router(config-isis)# address-family ipv4 unicast
RP/0//CPU0:router(config-isis-af)# maximum-redistributed-prefixes 5000 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 | <i>seconds</i> | 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 | <i>seconds</i> : 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 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.

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/RP0/CPU0:router(config)# router isis isp
RP/0/RP0/CPU0:router(config-isis)# max-lsp-lifetime 11000
```

| | | |
|-------------------------|---------------------------------------------------|--------------------------------|
| Related Commands | Command | Description |
| | lsp-refresh-interval, on page 584 | Sets the LSP refresh interval. |

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 | <i>number</i> 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 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.

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/RP0/CPU0:router(config)# router isis isp
RP/0/RP0/CPU0:router(config-isis)# interface GigabitEthernet 0/1/0/0
RP/0/RP0/CPU0:router(config-isis-if)# mesh-group 10
RP/0/RP0/CPU0:router(config-isis-if)# exit
RP/0/RP0/CPU0:router(config-isis)# interface GigabitEthernet 0/1/0/1
RP/0/RP0/CPU0:router(config-isis-if)# mesh-group 10
RP/0/RP0/CPU0:router(config-isis-if)# exit
RP/0/RP0/CPU0:router(config-isis)# interface GigabitEthernet 0/2/0/0
RP/0/RP0/CPU0:router(config-isis-if)# mesh-group 11
RP/0/RP0/CPU0:router(config-isis-if)# exit
RP/0/RP0/CPU0:router(config-isis)# interface GigabitEthernet 0/2/0/1
RP/0/RP0/CPU0:router(config-isis-if)# mesh-group 11
RP/0/RP0/CPU0:router(config-isis-if)# exit
RP/0/RP0/CPU0:router(config-isis)# interface GigabitEthernet 0/3/0/1
RP/0/RP0/CPU0:router(config-isis-if)# mesh-group 12
RP/0/RP0/CPU0:router(config-isis-if)# exit
RP/0/RP0/CPU0:router(config-isis)# interface GigabitEthernet 0/3/0/0
RP/0/RP0/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} [level {1 | 2}]
no metric [{default-metric | maximum} [level {1 | 2}]]
```

| Syntax Description | |
|--------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>default-metric</i> | 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). |
| 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 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.

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.

| 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:

```
RP/0/RP0/CPU0:router(config)# router isis isp
RP/0/RP0/CPU0:router(config-isis)# interface GigabitEthernet /1/0/1
RP/0/RP0/CPU0:router(config-isis-if)# address-family ipv4 unicast
RP/0/RP0/CPU0: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:

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router isis isp
RP/0/RP0/CPU0:router(config-isis)# address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-isis-af)# metric 15 level 2
```

Related Commands

| Command | Description |
|------------------------------------------------------|-------------------------------------------------------------------------------------------|
| metric-style narrow, on page 593 | Configures a router running IS-IS so that it generates and accepts old-style TLV objects. |
| metric-style transition, on page 595 | Configures the software to generate and accept both old-style and new-style TLV objects. |
| metric-style wide, on page 597 | 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 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.

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/RP0/CPU0:router(config)# router isis isp
RP/0/RP0/CPU0:router(config-isis)# address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-isis-af)# metric-style narrow level 1
```

Related Commands

| Command | Description |
|------------------------------------------------------|--------------------------------------------------------------------------------------|
| metric-style transition, on page 595 | Configures a router to generate and accept both old-style and new-style TLV objects. |
| metric-style wide, on page 597 | 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 | <p>transition Instructs the router to generate and accept both old-style and new-style TLV objects.</p> <p>level { 1 2 } (Optional) Specifies routing Level 1 or Level 2 independently.</p> | | | | |
|---------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------|--------------|---------------|------------------------------|
| Command Default | <p>Old-style TLVs are generated, if this command is not configured.</p> <p>Both Level 1 and Level 2 are configured if no level is specified.</p> | | | | |
| Command Modes | Address family configuration | | | | |
| Command History | <table border="1"> <thead> <tr> <th>Release</th> <th>Modification</th> </tr> </thead> <tbody> <tr> <td>Release 5.0.0</td> <td>This command was introduced.</td> </tr> </tbody> </table> | Release | Modification | Release 5.0.0 | This command was introduced. |
| Release | Modification | | | | |
| Release 5.0.0 | This command was introduced. | | | | |
| Usage Guidelines | <p>To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.</p> <p>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.</p> | | | | |
| Task ID | <table border="1"> <thead> <tr> <th>Task ID</th> <th>Operations</th> </tr> </thead> <tbody> <tr> <td>isis</td> <td>read, write</td> </tr> </tbody> </table> | Task ID | Operations | isis | read, write |
| 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/RP0/CPU0:router(config)# router isis isp
RP/0/RP0/CPU0:router(config-isis)# address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-isis-af)# metric-style transition level 2
```

Related Commands

| Command | Description |
|--------------------------------------------------|------------------------------------------------------------------------|
| metric-style narrow, on page 593 | Configures a router to generate and accept only old-style TLV objects. |
| metric-style wide, on page 597 | 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 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.

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/RP0/CPU0:router(config)# router isis isp  
RP/0/RP0/CPU0:router(config-isis)# address-family ipv4 unicast  
RP/0//CPU0:router(config-isis-af)# metric-style wide level 1
```

Related Commands

| Command | Description |
|--------------------------------------------------|------------------------------------------------------------------------|
| metric-style narrow, on page 593 | Configures a router to generate and accept only old-style TLV objects. |

microloop avoidance

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 <i>delay</i> | (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
```

microloop avoidance segment-routing

To enable the segment routing microloop avoidance and set the Routing Information Base (RIB) update delay value, use the **microloop avoidance** command. To disable segment routing microloop avoidance, use the **no** form of this command. Microloop avoidance segment-routing gets triggered by following events:

- link down
- link up
- change in link metrics
- overload bit set on node
- overload bit cleared on node

microloop avoidance segment-routing

Command Default

Disabled.

Command Modes

IPv4 address family configuration

Router 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.

The Segment Routing Microloop Avoidance feature detects if microloops are possible following a topology change. The node configures IS-IS prefixes with an explicit path by using a list of Segment Routing (SR) segments. The list of SR segments forces the traffic along the new path regardless whether nodes along the path already converged or not. This process eliminates the microloops. After the RIB update delay timer expires, the explicit list of SR segments is removed from the IS-IS prefixes.

Links or nodes that are not participating in the SPT (shortest path tree) of the given IS-IS level do not trigger the microloop avoidance.

Task ID

| Task ID | Operation |
|---------|-----------|
| ospf | read, |
| isis | write |

Example

This example shows how to enable Segment Routing Microloop Avoidance for IS-IS:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router isis 1
RP/0/RSP0/CPU0:router(config-isis)# address-family ipv4 unicast
RP/0/RSP0/CPU0:router(config-isis-af)# microloop avoidance segment-routing
RP/0/RSP0/CPU0:router(config-isis-af)# microloop avoidance rib-update-delay 3000
```

The `show isis protocol` command shows the configured uloop settings, for example, the topologies supported by IS-IS.

```
Router# show isis protocol
```

```
IPv4 Unicast
  Level-2
    Metric style (generate/accept): Wide/Wide
    Metric: 10
    Microloop avoidance: Enabled
      Configuration: Type: Segment routing, RIB update delay: 3000 msec
      State: Active, Duration: 2146 ms, Event Link down, Near: enxrr6.00 Far: enxrr5.00
```

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 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.

This command can be used to protect a router against the possible instability of its neighbor's LSPs.

The command parameters are similar 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/RP0/CPU0:router(config)# router isis isp  
RP/0/RP0/CPU0:router(config)# router isis isp min-lsp-arrivaltime  
RP/0/RP0/CPU0:router(config)# router isis 1 min- lsp-arrivaltime initial-wait  
RP/0/RP0/CPU0:router(config)#router isis 1 min-lsp-arrivaltime maximum-wait  
RP/0/RP0/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 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.

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/RP0/CPU0:router(config)# router isis isp
RP/0/RP0/CPU0:router(config-isis)# address-family ipv4 unicast
RP/0/RP0/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 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. |
|-------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|



| | |
|-------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 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/RP0/CPU0:router(config)# router isis isp
RP/0/RP0/CPU0:router(config-isis)# interface GigabitEthernet 0/3/0/0
RP/0/RP0/CPU0:router(config-isis-if)# address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-isis-if-af)# mpls ldp sync
```

| Related Commands | Command | Description |
|-------------------------|--------------------------------------------------|-------------------------------------------------|
| | show isis interface, on page 665 | 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 5.0.0 | This command was introduced. |
| Usage Guidelines | <p>To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.</p> <p>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.</p> | |
| 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/RP0/CPU0:router(config)# router isis isp
RP/0/RP0/CPU0:router(config-isis)# address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-isis-af)# mpls traffic-eng level-1
```

Related Commands

| Command | Description |
|-----------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------|
| mpls traffic-eng router-id (IS-IS), on page 611 | Specifies that the traffic engineering router identifier for the node is the IP address associated with a given interface. |

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 XR Config 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 XR Config

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | 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.

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/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# mpls traffic-eng path-selection ignore overload
```

The following example shows how to deactivate IS-IS overload bit avoidance:

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# no mpls traffic-eng path-selection ignore overload
```

Related Commands

| Command | Description |
|-----------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------|
| set-overload-bit, on page 640 | 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 | | |
|--------------------------|-------------|------------------------------------------------------------------------------------------------------------|
| <i>ip-address</i> | | IP address in four-part, dotted-decimal notation. |
| <i>type</i> | | Interface type. For more information, use the question mark (?) online help function. |
| <i>interface-path-id</i> | | 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 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.

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/RP0/CPU0:router(config)# router isis isp  
RP/0/RP0/CPU0:router(config-isis)# address-family ipv4 unicast  
RP/0/RP0/CPU0:router(config-isis-af)# mpls traffic-eng router-id Loopback0
```

Related Commands

| Command | Description |
|-------------------------------------------------------|----------------------------------------------------------------------------------------------------|
| mpls traffic-eng (IS-IS), on page 607 | Turns on flooding of MPLS traffic engineering link information in the indicated IGP level or area. |

mpls traffic-eng srlg (IS-IS)

To configure Shared Risk Link Group (SRLG) membership of a link, use the **mpls traffic-eng srlg** command in XR Config mode. To disable this feature, use the **no** form of this command.

```
mpls traffic-eng [type interface-path-id] srlg number
no mpls traffic-eng [type interface-path-id] srlg number
```

| Syntax Description | |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| number | SRLG membership number. Range is 0 to 4294967295. |
| 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 Shared Risk Link Group memberships are not configured.

Command Modes XR Config

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | 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 | Operations |
|---------|---------|----------------|
| | isis | read, write |
| | mpls-te | read, write |

Examples

The following example shows how to configure an SRLG with 10 member links:

```
RP/0/RP0/CPU0:router (config) #mpls traffic-eng interface gigabitEthernet 0/1/0/1
RP/0/RP0/CPU0:router (config-mpls-te-if) #srlg 10
```

Related Commands

| Command | Description |
|-------------------------------------------------------|----------------------------------------------------------------------------------------------------|
| mpls traffic-eng (IS-IS), on page 607 | 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 5.0.0 | 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/RP0/CPU0:router(config)# router isis isp  
RP/0/RP0/CPU0:router(config-isis)# net 47.0004.004d.0001.0001.0c11.1110.00
```

Related Commands

| Command | Description |
|------------------------------------------------------------|----------------------------------------------------------------------------|
| log adjacency changes (IS-IS), on page 572 | Configures the routing level for an instance of the IS-IS routing process. |
| router isis, on page 638 | 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 nsf** 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 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.

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/RP0/CPU0:router(config)# router isis isp
RP/0/RP0/CPU0:router(config-isis)# nsf cisco
```

Related Commands

| Command | Description |
|----------------------------------------------------|---------------------------------------------------------------------------------|
| nsf interface-expires, on page 619 | Configures the number of resends of an acknowledged NSF-restart acknowledgment. |

| Command | Description |
|---------------------------------------------------|--------------------------------------------------------------------------------------------------|
| nsf interface-timer, on page 621 | Configures the time interval after which an unacknowledged IETF NSF restart attempt is repeated. |
| nsf lifetime (IS-IS), on page 623 | 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 | <i>number</i> Number of resends. Range is 1 to 3. |
|---------------------------|---------------------------------------------------|

| | |
|------------------------|---------------------------|
| Command Default | <i>number</i> : 3 resends |
|------------------------|---------------------------|

| | |
|----------------------|----------------------|
| Command Modes | Router configuration |
|----------------------|----------------------|

| | | |
|------------------------|----------------|------------------------------|
| Command History | Release | Modification |
| | 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.

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/RP0/CPU0:router(config)# router isis isp
RP/0/RP0/CPU0:router(config-isis)# nsf ietf
RP/0/RP0/CPU0:router(config-isis)# nsf interface-expires 1
```

Related Commands

| Command | Description |
|--------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------|
| hello-multiplier, on page 549 | 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 621 | 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 | <i>seconds</i> NSF restart time interval (in seconds). Range is 3 to 20 seconds. |
|---------------------------|----------------------------------------------------------------------------------|

| | |
|------------------------|-----------------------------|
| Command Default | <i>seconds</i> : 10 seconds |
|------------------------|-----------------------------|

| | |
|----------------------|----------------------|
| Command Modes | Router configuration |
|----------------------|----------------------|

| Command History | Release | Modification |
|------------------------|----------------|------------------------------|
| | 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. |
|-------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

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/RP0/CPU0:router(config)# router isis isp
RP/0/RP0/CPU0:router(config-isis)# nsf ietf
RP/0/RP0/CPU0:router(config-isis)# nsf interface-timer 5
```

Related Commands

| Command | Description |
|-----------------------------------------------------|---------------------------------------------------------------------------------|
| nsf interface-expires, on page 619 | Configures the number of resends of an acknowledged NSF-restart acknowledgment. |
| hello-interval (IS-IS), on page 547 | 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 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.

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/RP0/CPU0:router(config)# router isis isp
RP/0/RP0/CPU0:router(config-isis)# nsf cisco
RP/0/RP0/CPU0:router(config-isis)# nsf lifetime 20
```

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 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 | 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/RP0/CPU0:router(config)# router isis isp
RP/0/RP0/CPU0:router(config-isis)# interface GigabitEthernet 0/1/0/1
RP/0/RP0/CPU0:router(config-isis-if)# passive
```

| Related Commands | Command | Description |
|------------------|-----------------------------------------|----------------------------------------------------------------------------------------------------------------------|
| | suppressed, on page 713 | 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 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.

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/RP0/CPU0:router(config)# router isis isp
RP/0/RP0/CPU0:router(config-isis)# interface TenGigE 0/6/0/0
RP/0/RP0/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

| | |
|--------------------------------------|----------------------------------------------------------------|
| <i>value</i> | 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 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.

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/RP0/CPU0:router(config)# router isis isp
```

```
RP/0/RP0/CPU0:router(config-isis)# interface TenGigE0/6/0/0  
RP/0/RP0/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 <i>route-policy-name</i> | | 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 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.

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/RP0/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/RP0/CPU0:router(config)# router isis isp
RP/0/RP0/CPU0:router(config-isis)# net 49.1234.2222.2222.00
RP/0/RP0/CPU0:router(config-isis)# address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-isis-af)# propagate level 2 into level 1 route-policy policy_a
```


Related Commands

| Command | Description |
|---------------------------------------------------|-------------------------------------------------------------------------------|
| redistribute (IS-IS), on page 630 | Redistributes routes from one routing domain into a specified IS-IS instance. |

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 | rib-metric-as-external | rib-metric-as-internal}] [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 | rib-metric-as-external | rib-metric-as-internal}] [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 | rib-metric-as-external | rib-metric-as-internal}] [route-policy route-policy-name]
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 | rib-metric-as-external
| rib-metric-as-internal}] [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 | rib-metric-as-external
| rib-metric-as-internal}] [route-policy route-policy-name]
no redistribute
```

Static Routes

```
redistribute static [{level-1 | level-2 | level-1-2}] [metric metric-value] [metric-type {{internal |
external | rib-metric-as-external} | {}]}] [route-policy route-policy-name]
no redistribute
```

| | | |
|---------------------------|-----------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax Description | <i>process-id</i> | <p>For the bgp keyword, an autonomous system number has the following ranges:</p> <ul style="list-style-type: none"> • 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. <p>For the isis keyword, an IS-IS instance identifier from which routes are to be redistributed.</p> <p>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.</p> <p>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.</p> |
| | 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 <i>metric-value</i> | (Optional) Specifies the metric used for the redistributed route. Range is 0 to 16777215. The <i>metric-value</i> 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 rib-metric-as-external rib-metric-as-internal } | <p>(Optional) Specifies the external link type associated with the route advertised into the ISIS routing domain. It can be one of four values:</p> <ul style="list-style-type: none"> • 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 • rib-metric-as-external–Use the rib-metric-as-external keyword to use RIB metric and set IS-IS external metric-type <p>Any route with an internal metric (however large the metric is) is preferred over a route with external metric (however small the metric is).</p> <p>Use the rib-metric-as-external and rib-metric-as-internal keywords to preserve RIB metrics when redistributing routes from another IS-IS router instance or another protocol.</p> |
| | route-policy <i>route-policy-name</i> | (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**] | **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 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.

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 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.



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 628](#) 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` readadvertises 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//CPU0:router(config)# router isis isp_A
RP/0//CPU0:router(config-isis)# net 49.1234.2222.2222.00
RP/0//CPU0:router(config-isis)# address-family ipv4 unicast
RP/0//CPU0:router(config-isis-af)# redistribute isis isp_B level-2
!
RP/0//CPU0:router(config)# router isis isp_B
RP/0//CPU0:router(config-isis)# is-type level 1
RP/0//CPU0:router(config-isis)# net 49.4567.2222.2222.00
RP/0//CPU0:router(config-isis)# address-family ipv4 unicast
```

Related Commands

| Command | Description |
|----------------------------------------------|------------------------------------------------------------|
| propagate level, on page 628 | Propagates routes from one IS-IS level into another level. |

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 | <i>seconds</i> | 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 | <i>seconds</i> : 5 seconds | |
| Command Modes | Interface configuration | |
| Command History | Release | Modification |
| | 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.

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 convergence. 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/RP0/CPU0:router(config)# router isis isp
RP/0/RP0/CPU0:router(config-isis)# interface GigabitEthernet 0/2/0/1
RP/0/RP0/CPU0:router(config-isis-if)# retransmit-interval 60
```

Related Commands

| Command | Description |
|-----------------------------------------------------------|--------------------------------------------------------------------------------------------------------|
| retransmit-throttle-interval, on page 636 | 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 | <i>milliseconds</i> | 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 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.

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/RP0/CPU0:router(config)# router isis isp
RP/0/RP0/CPU0:router(config-isis)# interface GigabitEthernet 0/2/0/1
RP/0/RP0/CPU0:router(config-isis-if)# retransmit-throttle-interval 300
```


Related Commands

| Command | Description |
|----------------------------------------------------------|----------------------------------------------------------------------------------------------------|
| lsp-gen-interval, on page 575 | Configures the minimum interval time between regenerating the same LSP. |
| retransmit-interval (IS-IS), on page 634 | 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 XR Config mode. To disable IS-IS routing, use the **no** form of this command.

```
router isis instance-id
no router isis instance-id
```

| | |
|---------------------------|-------------------------------------------------------------------------------------|
| Syntax Description | <i>instance-id</i> Name of the routing process. Maximum number of characters is 40. |
|---------------------------|-------------------------------------------------------------------------------------|

| | |
|------------------------|-------------------------------------------|
| Command Default | An IS-IS routing protocol is not enabled. |
|------------------------|-------------------------------------------|

| | |
|----------------------|-----------|
| Command Modes | XR Config |
|----------------------|-----------|

| Command History | Release | Modification |
|------------------------|----------------|------------------------------|
| | Release 5.0.0 | 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/RP0/CPU0:router(config)# router isis isp
RP/0/RP0/CPU0:router(config-isis)# net 49.0001.0000.0001.00
```

Related Commands

| Command | Description |
|----------------------------------|--------------------------------------------------|
| net, on page 615 | Configures an IS-IS NET for the routing process. |

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/RP0/CPU0:router(config)# router isis isp
RP/0/RP0/CPU0:router(config-isis)# address-family ipv4 multicast
RP/0/RP0/CPU0:router(config-isis-af)# route source first-hop
```

| Related Commands | Command | Description |
|------------------|----------------------------------|--------------------------------------------------|
| | net, on page 615 | Configures an IS-IS NET for the routing process. |

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. |
| <i>delay</i> | (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 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.

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/RP0/CPU0:router(config)# router isis isp  
RP/0/RP0/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 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.

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 630](#) 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/RP0/CPU0:router(config)# router isis 1
RP/0/RP0/CPU0:router(config-isis)# net 49.0001.0001.0001.0001.00
```

```

RP/0/RP0/CPU0:router(config-isis)# address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-isis-af)# redistribute isis 2 level 2
!
RP/0/RP0/CPU0:router(config-isis-af)# interface GigabitEthernet 0/3/0/0
RP/0/RP0/CPU0:router(config-isis-af-if)# address-family ipv4 unicast
!
!
RP/0/RP0/CPU0:router(config)# router isis 2
RP/0/RP0/CPU0:router(config-isis)# is-type level-1
RP/0/RP0/CPU0:router(config-isis)# net 49.0002.0001.0001.0002.00
RP/0/RP0/CPU0:router(config-isis)# address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-isis-af)# -attachedbit send always-
!
RP/0/RP0/CPU0:router(config-isis-af)# interface GigabitEthernet 0/1/0/0
RP/0/RP0/CPU0:router(config-isis-af-if)# address-family ipv4 unicast

```

Related Commands

| Command | Description |
|---------------------------------------------------|---------------------------------------------------------------------|
| redistribute (IS-IS), on page 630 | Redistributes routes from one IS-IS instance into another instance. |
| single-topology, on page 705 | 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

XR EXEC

Command History

Release

Release 5.0.0

Modification

This command was introduced.

Usage Guidelines

To use this command, you must be in a user group associated with a task 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. From Release 6.6.1, status of incremental shortest path first (iSPF) computation (enabled or not) is not listed.

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
      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 47: 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). From Release 6.6.1, status of incremental shortest path first (iSPF) computation (enabled or not) is not listed. |
| 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 XR EXEC mode.

show isis [**instance** *instance-id*] **adjacency** [**level** {**1** | **2**}] [*type interface-path-id*] [**detail**] [**systemid** *system-id*]

| Syntax Description | |
|--------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| instance <i>instance-id</i> | (Optional) Displays the IS-IS adjacencies for the specified IS-IS instance only. <ul style="list-style-type: none"> The <i>instance-id</i> 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. |
| <i>type</i> | Interface type. For more information, use the question mark (?) online help function. |
| <i>interface-path-id</i> | Physical interface or virtual interface. <p>Note Use the show interfaces command to see a list of all interfaces currently configured on the router.</p> <p>For more information about the syntax for the router, use the question mark (?) online help function.</p> |
| detail | (Optional) Displays neighbor IP addresses and active topologies. |
| systemid <i>system-id</i> | (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 XR EXEC

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | 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 | Operations |
|---------|---------|------------|
| | isis | read |

Examples The following is sample output from the **show isis adjacency** command:

```
RP/0/RP0/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 48: 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 <ul style="list-style-type: none"> • 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 679 | 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 XR 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 <i>number</i> of entries. Range is 1 to 100. | | | | |
| | first number (Optional) Specifies that the output is restricted to the first <i>number</i> of entries. Range is 1 to 100. | | | | |
| Command Default | No default behavior or values | | | | |
| Command Modes | XR EXEC | | | | |
| Command History | <table border="1"> <thead> <tr> <th>Release</th> <th>Modification</th> </tr> </thead> <tbody> <tr> <td>Release 5.0.0</td> <td>This command was introduced.</td> </tr> </tbody> </table> | Release | Modification | Release 5.0.0 | This command was introduced. |
| Release | Modification | | | | |
| 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 | <table border="1"> <thead> <tr> <th>Task ID</th> <th>Operations</th> </tr> </thead> <tbody> <tr> <td>isis</td> <td>read</td> </tr> </tbody> </table> | Task ID | Operations | isis | read |
| Task ID | Operations | | | | |
| isis | read | | | | |
| Examples | The following is sample output from the show isis adjacency-log command: | | | | |

```
RP/0/RP0/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 49: 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 XR EXEC mode.

show isis [**instance** *instance-id*] **checkpoint adjacency**

| | |
|---------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax Description | <p>instance <i>instance-id</i> (Optional) Displays the IS-IS checkpoint adjacencies for the specified IS-IS instance only.</p> <ul style="list-style-type: none"> The <i>instance-id</i> 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 | XR EXEC |
|----------------------|---------|

| Command History | Release | Modification |
|------------------------|----------------|------------------------------|
| | Release 5.0.0 | This command was introduced. |

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

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/RP0/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 50: 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 646 | Displays IS-IS adjacencies. |
| show isis checkpoint lsp, on page 654 | 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 XR EXEC mode.

show isis checkpoint interface

This command has no keywords or arguments.

Command Default No default behavior or values

Command Modes EXXR EXECEC

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | 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 | Operations |
|---------|---------|------------|
| | isis | read |

Examples

The following is sample output from the **show isis checkpoint interface** command:

```
RP/0/RP0/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 51: 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 XR EXEC mode.

show isis [**instance** *instance-id*] **checkpoint lsp**

| | | |
|---------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------|
| Syntax Description | instance <i>instance-id</i> (Optional) Displays the IS-IS checkpoint LSPs for the specified instance only. <ul style="list-style-type: none"> The <i>instance-id</i> 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 | XR EXEC | |
| Command History | Release | Modification |
| | Release 5.0.0 | This command was introduced. |
| Usage Guidelines | <p>To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.</p> <p>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.</p> | |
| Task ID | Task ID | Operations |
| | isis | read |
| Examples | The following is sample output from the show isis checkpoint lsp command: | |

```
RP/0/RP0/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 52: show isis checkpoint lsp Field Descriptions

| Field | Description |
|----------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Level | Routers with Level 1 or Level 2 adjacency configured. |
| LSPID | <p>LSP identifier. The first six octets form the system ID of the router that originated the LSP.</p> <p>The next octet is the pseudonode ID. When this byte is , 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.</p> <p>For each LAN, the designated router for that LAN creates and floods a pseudonode LSP, describing all systems attached to that LAN.</p> <p>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.</p> |
| Chkpt ID | Unique ID issued to the checkpoint at its creation. |

Related Commands

| Command | Description |
|-------------------------------------------------------------|---------------------------------------------------|
| show isis checkpoint adjacency, on page 650 | Displays the IS-IS checkpoint adjacency database. |
| show isis database, on page 656 | 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 XR EXEC mode.

```
show isis [instance instance-id] database [level {1 | 2}] [update] [summary] [detail] [verbose]
[{*lsp-id}]
```

Syntax Description

| | |
|--------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| instance <i>instance-id</i> | (Optional) Displays the IS-IS LSP database for the specified instance only. <ul style="list-style-type: none"> The <i>instance-id</i> 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. |
| * <i>lsp-id</i> | (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

XR EXEC

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 5.0.0 | 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 |

Examples

The following sample output shows the remaining lifetime value that is received in LSP database. The received value is shown in the output under the **Rcvd** field.

The following is sample output from the **show isis database** command with the **summary** keyword:

```
RP/0/RP0/CPU0:router# show isis database summary

IS-IS 10 Database Summary for all LSPs
                                Active          Purged          All
                                L1    L2    Total    L1    L2    Total    L1    L2    Total
-----
Fragment 0 Counts
  Router LSPs:      1     1     2     0     0     0     1     1     2
  Pseudo-node LSPs: 0     0     0     0     0     0     0     0     0
  All LSPs:        1     1     2     0     0     0     1     1     2
Per Topology
  IPv4 Unicast
  ATT bit set LSPs: 0     0     0     0     0     0     0     0     0
  OVL bit set LSPs: 0     0     0     0     0     0     0     0     0
All Fragment Counts
  Router LSPs:      1     1     2     0     0     0     1     1     2
  Pseudo-node LSPs: 0     0     0     0     0     0     0     0     0
  All LSPs:        1     1     2     0     0     0     1     1     2
```

This table describes the significant fields shown in the display.

Table 53: show isis database summary Field Descriptions

| Field | Description |
|-------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Router LSPs | Active, purged, and total LSPs associated with routers. |
| Pseudo-node LSPs: | Active, purged, and total LSPs associated with pseudonodes. |
| All LSPs: | Total active and purged LSPs. |
| ATT bit set LSPs | Attach bit (ATT). Indicates that the router is also a Level 2 router, and it can reach other areas. Level 1-only routers and Level 1-2 routers that have lost connection to other Level 2 routers use the Attach bit to find the closest Level 2 router. They point to a default route to the closest Level 2 router. |
| OVL bit set LSPs | Overload bit. Indicates 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. |

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 XR 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 <i>number</i> of entries. Range is 1 to 1000. |
| | first number (Optional) Specifies that the output be restricted to the first <i>number</i> of entries. Range is 1 to 1000. |

Command Default Both Level 1 and Level 2 are configured if no level is specified.

Command Modes XR EXEC

| | | |
|------------------------|----------------|------------------------------|
| Command History | Release | Modification |
| | 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 | Operations |
| | isis | read |

Examples

The following is sample output from the **show isis database-log** command:

```
RP/0/RP0/CPU0:router# show isis database-log

IS-IS 10 Level 1 Link State Database Log

```

| WHEN | LSPID | Op | Seq Num | Holdtime | OL | Old LSP | 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 54: show isis database-log Field Descriptions

| Field | Description |
|----------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| WHEN | Elapsed time (in hh:mm:ss) since the event was logged. |
| LSPID | <p>LSP identifier. The first six octets form the system ID of the router that originated the LSP.</p> <p>The next octet is the pseudonode ID. When this byte is , 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.</p> <p>For each LAN, the designated router for that LAN creates and floods a pseudonode LSP, describing all systems attached to that LAN.</p> <p>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.</p> |
| 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 656 | 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 XR EXEC mode.

show isis fast-reroute

A.B.C.D/length | **detail** | **summary** | **sr-only**

| | | |
|---------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------|
| Syntax Description | <i>A.B.C.D/length</i> 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 5.0.0 | 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/RP0/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/RP0/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/RP0/CPU0:router#show isis fast-reroute summary
IS-IS fr IPv4 Unicast FRR summary
```

| | | Critical Priority | High Priority | Medium Priority | Low Priority | Total |
|--------------------------|-------|----------------------|------------------|--------------------|-----------------|-------|
| 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/RP0/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 XR EXEC mode.

show isis [**instance** *instance-id*] **hostname**

| Syntax Description | <p>instance <i>instance-id</i> (Optional) Displays the IS-IS router name-to-system ID mapping table for the specified IS-IS instance only.</p> <p>The <i>instance-id</i> argument is the instance identifier (alphanumeric) defined by the router isis command.</p> | | | | | | |
|---------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------|--------------|---------------|------------------------------|------|------|
| 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 | XR EXEC | | | | | | |
| Command History | <table border="1"> <thead> <tr> <th>Release</th> <th>Modification</th> </tr> </thead> <tbody> <tr> <td>Release 5.0.0</td> <td>This command was introduced.</td> </tr> </tbody> </table> | Release | Modification | Release 5.0.0 | This command was introduced. | | |
| Release | Modification | | | | | | |
| Release 5.0.0 | This command was introduced. | | | | | | |
| Usage Guidelines | <p>To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.</p> <p>The show isis hostname command does not display entries if the dynamic hostnames are disabled.</p> | | | | | | |
| Task ID | <table border="1"> <thead> <tr> <th>Task ID</th> <th>Task</th> <th>Operations</th> </tr> </thead> <tbody> <tr> <td></td> <td>isis</td> <td>read</td> </tr> </tbody> </table> | Task ID | Task | Operations | | isis | read |
| Task ID | Task | Operations | | | | | |
| | isis | read | | | | | |
| Examples | <p>The following is sample output from the show isis hostname command with the instance and <i>instance-id</i> values specified:</p> <pre>RP/0/RP0/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</pre> | | | | | | |

This table describes the significant fields shown in the display.

Table 55: 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 559 | 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 XR 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 | XR EXEC | |
| Command History | Release | Modification |
| | 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 | Operations |
| | isis | read |

Examples

The following is sample output from the **show isis interface** command:

```
RP/0/RP0/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 Rermit 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 56: 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 Reremit 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 IIIH 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. |
| 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 <ul style="list-style-type: none"> • 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 57: 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 XR EXEC mode.

```
show isis [instance instance-id] lsp-log [level {1 | 2}] [{last number | first number}]
```

| Syntax Description | |
|------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| instance <i>instance-id</i> | (Optional) Displays the LSP log information for the specified IS-IS instance only. <ul style="list-style-type: none"> The <i>instance-id</i> 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 <i>number</i> | (Optional) Specifies that the output be restricted to the last <i>number</i> of entries. Range is 1 to 20. |
| first <i>number</i> | (Optional) Specifies that the output be restricted to the first <i>number</i> 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 XR EXEC

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | 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 | 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/RP0/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 58: 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 | <p>A list of all reasons that triggered an LSP rebuild. The triggers are</p> <ul style="list-style-type: none"> • 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 XR EXEC mode.

```
show isis [instance instance-id] mesh-group
```

| Syntax Description | <p>instance <i>instance-id</i> (Optional) Displays the mesh group information for the specified IS-IS instance only.</p> <ul style="list-style-type: none"> The <i>instance-id</i> 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 | XR EXEC | | | | |
| Command History | <table border="1"> <thead> <tr> <th>Release</th> <th>Modification</th> </tr> </thead> <tbody> <tr> <td>Release 5.0.0</td> <td>This command was introduced.</td> </tr> </tbody> </table> | Release | Modification | Release 5.0.0 | This command was introduced. |
| Release | Modification | | | | |
| 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 | <table border="1"> <thead> <tr> <th>Task ID</th> <th>Operations</th> </tr> </thead> <tbody> <tr> <td>isis</td> <td>read</td> </tr> </tbody> </table> | Task ID | Operations | isis | read |
| 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/RP0/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 59: 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 XR EXEC mode.

```
show isis [instance instance-id] mpls traffic-eng adjacency-log [{last number | first number}]
```

| Syntax Description | |
|------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| instance <i>instance-id</i> | (Optional) Displays the MPLS TE adjacency changes for the specified IS-IS instance only. <ul style="list-style-type: none"> The <i>instance-id</i> argument is the instance identifier (alphanumeric) defined by the router isis command. |
| last <i>number</i> | (Optional) Specifies that the output is restricted to last <i>number</i> of entries. Range is 1 to 20. |
| first <i>number</i> | (Optional) Specifies that the output is restricted to first <i>number</i> 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 XR EXEC

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | 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.

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/RP0/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 60: 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 675 | 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 XR EXEC mode.

```
show isis [instance instance-id] mpls traffic-eng advertisements
```

| | |
|---------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax Description | instance <i>instance-id</i> (Optional) Displays the latest flooded record from MPLS TE for the specified IS-IS instance only. <ul style="list-style-type: none"> The <i>instance-id</i> 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 | XR EXEC |
|----------------------|---------|

| Command History | Release | Modification |
|------------------------|----------------|------------------------------|
| | 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. |
|-------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

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/RP0/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 61: 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 673 | Displays a log of MPLS TE adjacency changes for IS-IS. |

show isis mpls traffic-eng tunnel

To display Multiprotocol Label Switching traffic engineering (MPLS TE) tunnel information for an Intermediate System-to-Intermediate System (IS-IS) instance, use the **show isis mpls traffic-eng tunnel** command in XR EXEC mode.

```
show isis [instance instance-id] mpls traffic-eng tunnel
```

| | |
|---------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax Description | <p>instance <i>instance-id</i> (Optional) Displays the MPLS TE tunnel information for the specified IS-IS instance only.</p> <ul style="list-style-type: none"> The <i>instance-id</i> argument is the instance identifier (alphanumeric) defined by the router isis command. |
|---------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

| | |
|------------------------|-----------------------------------------------------------------------------------------------|
| Command Default | No instance ID specified displays the MPLS TE tunnel information for all the IS-IS instances. |
|------------------------|-----------------------------------------------------------------------------------------------|

| | |
|----------------------|---------|
| Command Modes | XR EXEC |
|----------------------|---------|

| Command History | <table border="1"> <thead> <tr> <th>Release</th> <th>Modification</th> </tr> </thead> <tbody> <tr> <td>Release 5.0.0</td> <td>This command was introduced.</td> </tr> </tbody> </table> | Release | Modification | Release 5.0.0 | This command was introduced. |
|------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------|--------------|---------------|------------------------------|
| Release | Modification | | | | |
| Release 5.0.0 | This command was introduced. | | | | |

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

Use the **show isis** command to find the current status of MPLS TE tunnels.

Tunnels are used in IS-IS next-hop calculations.

| Task ID | <table border="1"> <thead> <tr> <th>Task ID</th> <th>Operations</th> </tr> </thead> <tbody> <tr> <td>isis</td> <td>read</td> </tr> </tbody> </table> | Task ID | Operations | isis | read |
|----------------|------------------------------------------------------------------------------------------------------------------------------------------------------|---------|------------|------|------|
| Task ID | Operations | | | | |
| isis | read | | | | |

Examples

The following is sample output from the **show isis mpls traffic-eng tunnel** command:

```
RP/0/RP0/CPU0:router# show isis mpls traffic-eng tunnel

ISIS isp Level-2 MPLS Traffic Engineering tunnels
System Id      Tunnel Name    Bandwidth    Nexthop      Metric    Mode
router-6       tu0           100000       172.18.1.6   0         Relative
```

This table describes the significant fields shown in the display.

Table 62: show isis mpls traffic-eng tunnel 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. |
| Tunnel Name | Name of the MPLS TE tunnel interface. |
| Bandwidth | MPLS TE-specified tunnel bandwidth of the tunnel. |
| Nexthop | MPLS TE destination IP address of the tunnel. |
| Metric | MPLS TE metric of the tunnel. |
| Mode | MPLS TE metric mode of the tunnel. It can be relative or absolute. |

show isis neighbors

To display information about Intermediate System-to-Intermediate System (IS-IS) neighbors, use the **show isis neighbors** command in XR EXEC mode.

```
show isis [instance instance-id] neighbors [{type interface-path-id | summary}] [detail] [systemid system-id]
```

| Syntax Description | |
|------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| instance <i>instance-id</i> | (Optional) Displays the IS-IS neighbor information for the specified IS-IS instance only. <ul style="list-style-type: none"> The <i>instance-id</i> 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. <p>Note Use the show interfaces command to see a list of all interfaces currently configured on the router.</p> <p>For more information about the syntax for the router, use the question mark (?) online help function.</p> |
| summary | (Optional) Displays neighbor status count for each level. |
| detail | (Optional) Displays additional details. |
| systemid <i>system-id</i> | (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 XR EXEC

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | 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 | 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//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' 'IPv6 Unicast'
  Uptime: 01:09:44
  IPFRR: LFA Neighbor: elise
         LFA IPv4 address: 10.100.1.2
         LFA Router address: 192.168.0.45
LFA Interface: Gi0/1/0/0.
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' 'IPv6 Unicast'
  Uptime: 01:09:33

Total neighbor count: 3
```

This table describes the significant fields shown in the display.

Table 63: 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/RP0/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 64: 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 646 | 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 XR EXEC mode.

show isis [**instance** *instance-id*] **protocol**

| Syntax Description | instance <i>instance-id</i> (Optional) Displays the IS-IS adjacencies for the specified IS-IS instance only. <ul style="list-style-type: none"> The <i>instance-id</i> 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 | XR EXEC | | | | |
| Command History | <table border="1"> <thead> <tr> <th>Release</th> <th>Modification</th> </tr> </thead> <tbody> <tr> <td>Release 5.0.0</td> <td>This command was introduced.</td> </tr> </tbody> </table> | Release | Modification | Release 5.0.0 | This command was introduced. |
| Release | Modification | | | | |
| 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 | 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
      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)
```

This table describes the significant fields shown in the display.

Table 65: 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: <ul style="list-style-type: none"> • Cisco Proprietary NSF Restart • IETF NSF Restart • Cold Restart |
| No protocols redistributed: | No redistributed protocol information exists to be displayed. |
| Distance: | Administrative distance for this protocol. |

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 XR EXEC mode.

```
show isis [instance instance-id] [{ipv4 | ipv6 | afi-all}] [{unicast | [topology {alltopo-name}] | safi-all}] route [{ip-address mask | ip-address/length [longer-prefixes]}] [summary] [backup] [detail] [sr-only]
```

| Syntax Description | |
|------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| instance <i>instance-id</i> | (Optional) Displays the IP reachability information for the specified IS-IS instance only. <ul style="list-style-type: none"> The <i>instance-id</i> 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. |
| topology | (Optional) Specifies IS-IS paths to intermediate systems. |
| all | (Optional) Specifies all topologies. |
| topology <i>topo-name</i> | (Optional) Specifies topology table information and name of the topology table. |
| safi-all | (Optional) Specifies all secondary address prefixes. |
| <i>ip-address</i> | (Optional) Network IP address about which routing information should be displayed. |
| <i>mask</i> | (Optional) Network mask specified in either of two ways: <ul style="list-style-type: none"> 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. |
| <i>/ length</i> | (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. |
| backup | (Optional) Displays backup information for this entry. |

| | |
|---------------|------------------------------------------------------|
| 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

XR EXEC

Command History

| Release | Modification |
|---------------|---------------------------------------|
| Release 5.0.0 | 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 |

Examples

The following is sample output from the **show isis route** command:

```
RP/0//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 66: 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//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 XR EXEC mode.

```
show isis [ instance instance-id ] [ [ { ipv4 | ipv6 | afi-all } ] [ { unicast | [ topology { all
topo-name } ] | safi-all } ] ] spf-log [ level { 1 | 2 } ] [ { fspf | prc | nhc } ] [ { detail | verbose |
plfrr | ppfrr } ] [ { last number | first number } ]
```

| Syntax Description | instance <i>instance-id</i> | (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. |
| | topology all <i>topo-name</i> | (Optional) Specifies topology table information for all topologies or for the specified topology table (<i>topo-name</i>). |
| | 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. |
| | fsfp | (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 <i>number</i> | (Optional) Specifies that the output is restricted to the last <i>number</i> of entries. Range is 1 to 210. |
| | first <i>number</i> | (Optional) Specifies that the output is restricted to the first <i>number</i> 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, and prc).

Command Modes

XR EXEC

show isis spf-log

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | 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 | 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 67: 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. |
| 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 <i>List of Triggers</i> . |
| 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 <i>List of Triggers</i> . |

This table lists triggers of a full SPF calculation.

Table 68: List of Triggers

| Trigger | Description |
|---------------|------------------------------------------------------------------------------------------------------------------------------------------|
| PERIODIC | Runs a full SPF calculation every 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. |
| 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. |
| 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. |

| Trigger | Description |
|-------------|-----------------------------------|
| PREFIXTLV | Changed Prefix TLV content. |
| AREAADDRTLV | Changed Area address TLV content. |
| IP ADDR TLV | 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

      ISIS 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          12a5.00-00 NEWADJ

      ISIS 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 69: 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 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 <i>List of Triggers</i> . |

| Field | Description |
|-------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 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 <i>List of Triggers</i> . |

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 70: show isis spf-log detail Field Descriptions

| Field | Description |
|-----------|--------------------------------------------|
| Level | IS-IS level of the router. |
| Timestamp | Time at which the SPF calculation started. |

| Field | Description |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Type | Type of route calculation. The possible types are incremental 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 <i>List of Triggers</i> . |
| 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 <i>List of Triggers</i> . |
| Delay | Two different delays exist: <ol style="list-style-type: none"> 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: <ol style="list-style-type: none"> 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: <ol style="list-style-type: none"> 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. |
| 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. |

| Field | Description |
|-------|----------------------------------------------------------|
| Total | Total number of nodes or prefixes at various priorities. |

Related Commands

| Command | Description |
|-------------------------------------------|------------------------------------------------------------------|
| spf-interval, on page 707 | 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 XR EXEC mode.

show isis [**instance** *instance-id*] **statistics** [*type interface-path-id*]

| Syntax Description | |
|------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| instance <i>instance-id</i> | (Optional) Displays the IS-IS traffic statistics for the specified IS-IS instance only. <ul style="list-style-type: none"> The <i>instance-id</i> argument is the instance identifier (alphanumeric) defined by the router isis command. |
| <i>type</i> | Interface type. For more information, use the question mark (?) online help function. |
| <i>interface-path-id</i> | Physical interface or virtual interface. <p>Note Use the show interfaces command to see a list of all interfaces currently configured on the router.</p> <p>For more information about the syntax for the router, use the question mark (?) online help function.</p> |

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 XR EXEC

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | 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.

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
Level-2:
  LSPs sourced (new/refresh): 4255/3332
  IPv4 Unicast
    SPF calculations      : 432
    ISPF calculations     : 0
    Next-hop Calculations : 8
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 71: 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 submenu. |
| 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 687 | Displays how often and why the router has run a full SPF calculation. |
| spf-interval, on page 707 | 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 XR EXEC mode.

```
show isis [instance instance-id] [[{ipv4 | ipv6 | afi-all}] [{unicast | [topology {all | topo-name}] | safi-all}]] | summary | level {1 | 2} [systemid system-id] [detail]
```

Syntax Description

| | |
|------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| instance <i>instance-id</i> | (Optional) Displays the IS-IS topology for the specified IS-IS instance only. <ul style="list-style-type: none"> The <i>instance-id</i> 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. |
| topology <i>topo-name</i> | (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. |
| systemid <i>system-id</i> | (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

XR EXEC

Command History

| Release | Modification |
|---------------|------------------------------|
| 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.

Use the **show isis topology** command to verify the presence and connectivity among all routers in all areas.

| Task ID | Task ID | Operations |
|---------|---------|------------|
| | isis | read |

Examples

The following is sample output from the **show isis topology** command:

```
RP/0/RP0/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 72: 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/RP0/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 73: 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 XR EXEC mode.

```
show protocols [{afi-all | ipv4 | ipv6}] [{allprotocol}]
```

| 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. |
| <i>protocol</i> | (Optional) Specifies a routing protocol. For the IPv4 address family, the options are: <ul style="list-style-type: none"> • bgp • isis • ospf • rip • eigrp For the IPv6 address family, the options are: <ul style="list-style-type: none"> • bgp • isis • ospfv3 |

Command Default If no address family is specified, the default is IPv4.

Command Modes XR EXEC

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | 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 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 |

| Task ID | Operations |
|---------|------------|
| 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
      Level-2
        Metric style (generate/accept): Narrow/Narrow
      Redistributing:
        static
        Distance: 115
    IPv6 Unicast
      Level-1
      Level-2
      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 74: 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. |
| 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. |

| Field | Description |
|-------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 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 593 , metric-style transition, on page 595 , or metric-style wide, on page 597 command. |
| 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 593 | Configures the IS-IS software to generate and accept old-style type, length, and value (TLV) objects. |
| metric-style transition, on page 595 | 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 597 | Configures the IS-IS software to generate and accept only new-style type, length, and value (TLV) objects. |
| redistribute (IS-IS), on page 630 | 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 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 | Operations |
|---------|---------|----------------|
| | isis | read, write |

Examples

The following example disables the IS-IS protocol on GigabitEthernet interface 0/1/0/1:

```
RP/0/RP0/CPU0:router(config)# router isis isp
RP/0/RP0/CPU0:router(config-isis)# interface GigabitEthernet0/1/0/1
RP/0/RP0/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 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.

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/RP0/CPU0:router(config)# router isis isp
RP/0/RP0/CPU0:router(config-isis)# net 49.0000.0000.0001.00
RP/0/RP0/CPU0:router(config-isis)# address-family ipv6 unicast
RP/0/RP0/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}
```

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | 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.

Examples

```
RP/0/RP0/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/RP0/CPU0:router (config) #snmp-server traps isis all

RP/0/RP0/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 <i>initial</i> | | Initial SPF calculation delay (in milliseconds) after a topology change. Range is 0 to 120000. |
| secondary-wait <i>secondary</i> | | Hold time between the first and second SPF calculations (in milliseconds). Range is 0 to 120000. |
| maximum-wait <i>maximum</i> | | 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 <i>initial</i> : 50 milliseconds | |
| secondary-wait <i>secondary</i> : 200 milliseconds | |
| maximum-wait <i>maximum</i> : 5000 milliseconds | |

| Command Modes | |
|---------------|------------------------------|
| | Address family configuration |

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | 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.

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/RP0/CPU0:router(config)# router isis isp  
RP/0/RP0/CPU0:router(config-isis)# address-family ipv4 unicast  
RP/0/RP0/CPU0:router(config-isis-af)# spf-interval initial-wait 10 maximum-wait 5000
```


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. |
| | <i>access-list-name</i> | Name of an access list. |
| | tag <i>tag</i> | Specifies a tag to indicate priority. The <i>tag</i> 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 5.0.0 | This command was introduced. |
| Usage Guidelines | <p>To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.</p> <p>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:</p> <p>Critical > High > Medium > Low</p> <p>The spf prefix-priority command supports prefix lists for the first three priorities. The unmatched prefixes are updated with low priority.</p> <p>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.</p> | |
| Task ID | Task ID | Operations |
| | isis | read, write |

Examples

The following example shows how to set the prefix priorities:

```
RP/0/RP0/CPU0:router(config)# ipv4 prefix-list isis-critical-acl
RP/0/RP0/CPU0:router(config-ipv4_pfx)# 10 permit 0.0.0.0/0 eq 32
!
RP/0/RP0/CPU0:router(config)# ipv4 prefix-list isis-med-acl
RP/0/RP0/CPU0:router(config-ipv4_pfx)# 10 permit 0.0.0.0/0 eq 29
!
RP/0/RP0/CPU0:router(config)# ipv4 prefix-list isis-high-acl
RP/0/RP0/CPU0:router(config-ipv4_pfx)# 10 permit 0.0.0.0/0 eq 30
!
RP/0/RP0/CPU0:router(config)# router isis ring
RP/0/RP0/CPU0:router(config-isis)# address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-isis-af)# spf prefix-priority critical isis-critical-acl
RP/0/RP0/CPU0:router(config-isis-af)# spf prefix-priority high isis-high-acl
RP/0/RP0/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 <i>address</i> argument must be in four-part, dotted-decimal notation. |
| | / <i>prefix-length</i> | 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 <i>ipv6-prefix</i> 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 <i>tag</i> | 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 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.

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**

| Task ID | Operations |
|---------|----------------|
| isis | read, write |

Examples

```
RP/0/RP0/CPU0:router(config)# router isis isp
RP/0/RP0/CPU0:router(config-isis)# address-family ipv6 unicast
RP/0/RP0/CPU0:router(config-isis-af)# redistribute 2 level-2
RP/0/RP0/CPU0:router(config-isis-af)# summary-prefix level-2
RP/0/RP0/CPU0:router(config-isis-af)# summary-prefix
```

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 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.

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/RP0/CPU0:router(config)# router isis isp
RP/0/RP0/CPU0:router(config-isis)# interface GigabitEthernet /1/0/1
RP/0/RP0/CPU0:router(config-isis-if)# suppressed
```

| Related Commands | Command | Description |
|------------------|----------------------------------------------|------------------------------------------|
| | passive (IS-IS), on page 624 | 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 | <i>tag</i> 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 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 | Operations |
|----------------|----------------|-------------------|
| | isis | read, write |

| | |
|-----------------|------------------------------------------------------------------------------|
| Examples | The following example shows how to associate and advertise an interface tag: |
|-----------------|------------------------------------------------------------------------------|

```
RP/0/RP0/CPU0:router(config)# router isis isp
RP/0/RP0/CPU0:router(config-isis)# interface GigabitEthernet 0/3/0/0
RP/0/RP0/CPU0:router(config-isis-if)# address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-isis-if-af)# tag 234
```

| Related Commands | Command | Description |
|-------------------------|----------------------------------------------------------|-------------------------------------------------------------------------------|
| | spf prefix-priority (IS-IS), on page 709 | Assigns a priority to an ISIS prefix for customizing the RIB update sequence. |

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 |
| | <i>max-trace-entries</i> | 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 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 |

Examples

The following example shows how to set the isis buffer size for severe traces to 1200:

```
RP/0/RP0/CPU0:router(config)#router isis isp
RP/0/RP0/CPU0:router(config-isis)#trace sever 1200
```




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* module in the *Routing Configuration Guide for Cisco NCS 6000 Series Routers*.

- [address-family \(OSPF\)](#), on page 720
- [adjacency stagger](#), on page 721
- [area \(OSPF\)](#), on page 723
- [authentication \(OSPF\)](#), on page 725
- [authentication-key \(OSPF\)](#), on page 728
- [auto-cost \(OSPF\)](#), on page 730
- [capability opaque disable](#), on page 732
- [clear ospf process](#), on page 733
- [clear ospf redistribution](#), on page 735
- [clear ospf routes](#), on page 736
- [clear ospf statistics](#), on page 737
- [clear ospf statistics interface](#), on page 739
- [cost \(OSPF\)](#), on page 740
- [cost-fallback \(OSPF\)](#), on page 742
- [database-filter all out \(OSPF\)](#), on page 744
- [dead-interval \(OSPF\)](#), on page 745
- [default-cost \(OSPF\)](#), on page 747
- [default-information originate \(OSPF\)](#), on page 749
- [default-metric \(OSPF\)](#), on page 751
- [demand-circuit \(OSPF\)](#), on page 753
- [distance \(OSPF\)](#), on page 755
- [distance ospf](#), on page 758
- [distribute-list](#), on page 760
- [fast-reroute \(OSPFv2\)](#), on page 762
- [fast-reroute per-link exclude interface](#), on page 764
- [fast-reroute per-prefix exclude interface \(OSPFv2\)](#), on page 766
- [fast-reroute per-prefix lfa-candidate \(OSPFv2\)](#), on page 767
- [fast-reroute per-prefix remote-lfa \(OSPFv2\)](#), on page 768
- [fast-reroute per-prefix use-candidate-only \(OSPFv2\)](#), on page 770

- flood-reduction (OSPF), on page 771
- hello-interval (OSPF), on page 773
- interface (OSPF), on page 775
- log adjacency changes (OSPF), on page 777
- loopback stub-network , on page 778
- max-lsa, on page 779
- max-metric, on page 781
- maximum interfaces (OSPF), on page 784
- maximum paths (OSPF), on page 785
- maximum redistributed-prefixes (OSPF), on page 787
- message-digest-key, on page 789
- mpls ldp auto-config (OSPF), on page 792
- mpls ldp sync (OSPF), on page 793
- mpls ldp sync-igp-shortcuts, on page 794
- mpls traffic-eng ldp-sync-update (OSPF), on page 796
- mtu-ignore (OSPF), on page 797
- multi-area-interface, on page 799
- neighbor (OSPF), on page 801
- neighbor database-filter all out, on page 803
- network (OSPF), on page 804
- nsf (OSPF), on page 806
- nsf flush-delay-time (OSPF), on page 808
- nsf interval (OSPF), on page 809
- nsf lifetime (OSPF), on page 810
- nssa (OSPF), on page 811
- ospf name-lookup, on page 813
- packet-size (OSPF), on page 814
- passive (OSPF), on page 816
- priority (OSPF), on page 818
- protocol shutdown, on page 820
- queue dispatch flush-lsa, on page 821
- queue dispatch incoming, on page 823
- queue dispatch rate-limited-lsa, on page 825
- queue dispatch spf-lsa-limit, on page 827
- queue limit, on page 828
- range (OSPF), on page 830
- redistribute (OSPF), on page 832
- retransmit-interval (OSPF), on page 837
- route-policy (OSPF), on page 839
- router-id (OSPF), on page 840
- router ospf, on page 842
- security ttl (OSPF), on page 844
- show ospf, on page 846
- show ospf border-routers, on page 849
- show ospf flood-list, on page 851
- show ospf interface, on page 853

- [show ospf mpls traffic-eng](#), on page 855
- [show ospf message-queue](#), on page 860
- [show ospf neighbor](#), on page 863
- [show ospf request-list](#), on page 870
- [show ospf retransmission-list](#), on page 872
- [show ospf routes](#), on page 874
- [show ospf sham-links](#), on page 878
- [show ospf statistics interface](#), on page 879
- [show ospf summary-prefix](#), on page 881
- [show ospf virtual-links](#), on page 883
- [show protocols \(OSPF\)](#), on page 885
- [snmp trap rate-limit \(OSPF\)](#), on page 888
- [spf prefix-priority \(OSPFv2\)](#), on page 889
- [stub \(OSPF\)](#), on page 891
- [summary-prefix \(OSPF\)](#), on page 893
- [timers lsa group-pacing](#), on page 895
- [timers lsa min-arrival](#), on page 896
- [timers lsa refresh](#), on page 897
- [timers throttle lsa all \(OSPF\)](#), on page 899
- [timers throttle spf \(OSPF\)](#), on page 902
- [transmit-delay \(OSPF\)](#), on page 904
- [ucmp \(OSPFv2\)](#), on page 906
- [ucmp delay-interval \(OSPFv2\)](#), on page 908
- [ucmp exclude interface \(OSPFv2\)](#), on page 909
- [virtual-link \(OSPF\)](#), on page 911

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 | | | | |
| Command History | <table border="1"> <thead> <tr> <th>Release</th> <th>Modification</th> </tr> </thead> <tbody> <tr> <td>Release 5.0.0</td> <td>This command was introduced.</td> </tr> </tbody> </table> | Release | Modification | Release 5.0.0 | This command was introduced. |
| Release | Modification | | | | |
| Release 5.0.0 | This command was introduced. | | | | |
| Usage Guidelines | <p>To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.</p> <p>OSPF version 2 automatically provides routing services for IPv4 unicast topologies, so this command is redundant.</p> | | | | |
| Task ID | <table border="1"> <thead> <tr> <th>Task ID</th> <th>Operations</th> </tr> </thead> <tbody> <tr> <td>ospf</td> <td>read, write</td> </tr> </tbody> </table> | Task ID | Operations | ospf | read, write |
| Task ID | Operations | | | | |
| ospf | read, write | | | | |
| Examples | <p>The following example shows how to configure the OSPF router process with IPv4 unicast address prefixes:</p> <pre>RP/0/RP0/CPU0:router(config)# router ospf 1 RP/0/RP0/CPU0:router(config-ospf)# address-family ipv4 unicast</pre> | | | | |

adjacency stagger

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. |
|--------------------|------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| | <i>initial-num-nbr</i> | 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. |
| | <i>max-num-nbr</i> | 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 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.

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/RP0/CPU0:router# configure  
RP/0/RP0/CPU0:router(config)# router ospf 1  
RP/0/RP0/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 | <i>area-id</i> Identifier of an OSPF area. The <i>area-id</i> 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 | | | | |
| Command History | <table border="1"> <thead> <tr> <th>Release</th> <th>Modification</th> </tr> </thead> <tbody> <tr> <td>Release 5.0.0</td> <td>This command was introduced.</td> </tr> </tbody> </table> | Release | Modification | Release 5.0.0 | This command was introduced. |
| Release | Modification | | | | |
| 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.

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/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router ospf 1
RP/0/RP0/CPU0:router(config-ospf)# area 0
```

```
RP/0/RP0/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 ] | [keychain keychain ] | null}]  
no authentication
```

| Syntax Description | |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------|
| message-digest | (Optional) Specifies that MD5 is used. |
| keychain <i>keychain</i> | (Optional) Specifies a keychain name. |
| keychain <i>keychain</i> | (Optional) Enables Keychain-based authentication. |
| 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 |
| | Multi-area interface configuration |
| | Sham-link configuration |

| Command History | Release | Modification |
|-----------------|---------------|-------------------------------------------------------------------------------------|
| | Release 5.0.0 | This command was introduced. |
| | Release 6.5.1 | The Keychain keyword was introduced to enable Keychain-based authentication. |

| Usage Guidelines | |
|------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| | To use this command, you must be in a user group associated with a task group that 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/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router ospf 201
RP/0/RP0/CPU0:router(config-ospf)# router-id 10.1.1.1
RP/0/RP0/CPU0:router(config-ospf)# area 0
RP/0/RP0/CPU0:router(config-ospf-ar)# authentication
RP/0/RP0/CPU0:router(config-ospf-ar)# interface GigabitEthernet 0/1/0/1
RP/0/RP0/CPU0:router(config-ospf-ar-if)# authentication-key mykey
RP/0/RP0/CPU0:router(config-ospf-ar-if)# exit
RP/0/RP0/CPU0:router(config-ospf)# area 1
RP/0/RP0/CPU0:router(config-ospf-ar)# authentication
RP/0/RP0/CPU0:router(config-ospf-ar)# interface GigabitEthernet 0/1/0/0
RP/0/RP0/CPU0:router(config-ospf-ar-if)# authentication-key mykey1
```

The following example shows how to configure use of an authentication keychain:

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router ospf 201
RP/0/RP0/CPU0:router(config-ospf)# router-id 10.1.1.1
RP/0/RP0/CPU0:router(config-ospf)# authentication message-digest keychain mykeychain
```

The following example shows how to configure keychain-based authentication:

```
RP/0/RP0/CPU0:router(config-ospf)#authentication keychain ?
WORD The keychain name
RP/0/RP0/CPU0:router(config-ospf)#authentication keychain test_chain ----- Router level authentication
```

```

RP/0/RP0/CPU0:router(config-ospf)#router-id 1.1.1.1
RP/0/RP0/CPU0:router(config-ospf)#address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-ospf)#area 1
RP/0/RP0/CPU0:router(config-ospf-ar)#authentication keychain test_chain -----□ Area level
authentication
RP/0/RP0/CPU0:router(config-ospf-ar)#interface TenGigE 0/0/0/1
RP/0/RP0/CPU0:router(config-ospf-ar-if)#authentication keychain test_chain -----□ Interface
level authentication
RP/0/RP0/CPU0:router(config-ospf-ar-if)#commit

```

Related Commands

| Command | Description |
|--------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------|
| authentication-key (OSPF), on page 728 | Assigns a password to be used by neighboring routers that are using the simple password authentication of OSPF. |
| message-digest-key, on page 789 | 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. |
| | <i>password</i> | Any contiguous string up to 8 characters in length that can be entered from the keyboard. For example, <i>mypswd2</i> . |
| 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 | |
| | Multi-area configuration | |
| | Sham-link configuration | |
| Command History | Release | Modification |
| | 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. | |
| | 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/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router ospf 201
RP/0/RP0/CPU0:router(config-ospf)# authentication-key yourpass
```

Related Commands

| Command | Description |
|----------------------------------------------------|--------------------------------|
| authentication (OSPF), on page 725 | 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 <i>mbps</i> | 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

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | 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.

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/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router ospf 1
RP/0/RP0/CPU0:router(config-ospf)# auto-cost reference-bandwidth 1000
```

Related Commands

| Command | Description |
|------------------------------------------|-------------------------------------------------------------------------------------|
| cost (OSPF), on page 740 | 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

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | 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.

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/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router (config)# router ospf 1
RP/0/RP0/CPU0:router (config-ospf)# capability opaque disable
```


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 XR EXEC mode.

```
clear ospf process
```

| | |
|---------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax Description | <i>process-name</i> (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. |
|---------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

| | |
|------------------------|------------------------------|
| Command Default | No default behavior or value |
|------------------------|------------------------------|

| | |
|----------------------|---------|
| Command Modes | XR EXEC |
|----------------------|---------|

| Command History | Release | Modification |
|------------------------|----------------|------------------------------|
| | 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.

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 840](#) command.

| Task ID | Task ID | Operations |
|----------------|----------------|-------------------|
| | ospf | read, write |

Examples

The following example shows how to reset all OSPF processes:

```
RP/0/RP0/CPU0:router# clear ospf process
```

The following example shows how to reset the OSPF 1 process:

```
RP/0/RP0/CPU0:router# clear ospf 1 process
```

Related Commands

| Command | Description |
|-----------------------------------------------|----------------------------------------------|
| router ospf, on page 842 | Configures an OSPF routing process. |
| router-id (OSPF), on page 840 | 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 XR EXEC mode.

clear ospf redistribution

| | |
|---------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax Description | <i>process-name</i> (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. |
|---------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

| | |
|------------------------|------------------------------|
| Command Default | No default behavior or value |
|------------------------|------------------------------|

| | |
|----------------------|---------|
| Command Modes | XR EXEC |
|----------------------|---------|

| | | |
|------------------------|----------------|------------------------------|
| Command History | Release | Modification |
| | 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.

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/RP0/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 XR EXEC mode.

clear ospf routes

| | |
|---------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax Description | <i>process-name</i> (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. |
|---------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

| | |
|------------------------|------------------------------|
| Command Default | No default behavior or value |
|------------------------|------------------------------|

| | |
|----------------------|---------|
| Command Modes | XR EXEC |
|----------------------|---------|

| Command History | Release | Modification |
|------------------------|----------------|------------------------------|
| | 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 | Operations ID |
|----------------|-------------|----------------------|
| | 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/RP0/CPU0:router# clear ospf routes
```

| Related Commands | Command | Description |
|-------------------------|------------------------------------------|-------------------------------------|
| | router ospf, on page 842 | 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 XR EXEC mode.

```
clear ospf statistics [neighbor [type interface-path-id] [ip-address]]
```

Syntax Description

| | |
|--------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>process-name</i> | (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. |
| neighbor | (Optional) Clears the state transition counters of the specified neighbor only. |
| <i>type</i> | (Optional) Interface type. For more information, use the question mark (?) online help function. |
| <i>interface-path-id</i> | (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. |
| <i>ip-address</i> | (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

XR EXEC

Command History

| Release | Modification |
|---------------|------------------------------|
| 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.

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/RP0/CPU0:router# clear ospf statistics neighbor POS 0/2/0/0
```

Related Commands

| Command | Description |
|------------------------------------------|-------------------------------------|
| router ospf, on page 842 | Configures an OSPF routing process. |

clear ospf statistics interface

To clear the Open Shortest Path First (OSPF) statistics per interface, use the **clear ospf statistics interface** command in XR EXEC mode.

```
clear ospf statistics interface type interface-path-id
```

| | | |
|---------------------------|--------------------------|------------------------------------------------------------------------------------------------------------|
| Syntax Description | <i>type</i> | Interface type. For more information, use the question mark (?) online help function. |
| | <i>interface-path-id</i> | 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 value.

Command Modes XR EXEC

| | | |
|------------------------|----------------|------------------------------|
| Command History | Release | Modification |
| | 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.

Use the **clear ospf statistics interface** 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 OSPF statistics for interface POS 0/21/0/0:

```
RP/0/RP0/CPU0:router# clear ospf statistics interface POS 0/21/0/0
```

| | | |
|-------------------------|----------------------------------------------------|--------------------------------------------------------------------------------------|
| Related Commands | Command | Description |
| | clear ospf statistics, on page 737 | Clears the Open Shortest Path First (OSPF) statistics of neighbor state transitions. |

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

Multi-area configuration

Sham-link configuration

Command History

Release

Release 5.0.0

Modification

This command was introduced.

Usage Guidelines

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

$108 / \text{bandwidth}$ (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/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router (config)# router ospf 1
RP/0/RP0/CPU0:router (config-ospf)# area 0
RP/0/RP0/CPU0:router (config-ospf-ar)# interface GigabitEthernet 0/1/0/1
RP/0/RP0/CPU0:router (config-ospf-ar-if)# cost 65
```

Related Commands

| Command | Description |
|-----------------------------------------------|------------------------------------------------------------------------------|
| auto-cost (OSPF), on page 730 | 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

| | | |
|------------------|------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>cost</i> | 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. |
| <i>bandwidth</i> | | 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 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.

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 for Packet-over-SONET/SDH (POS):

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)#router ospf 100
RP/0/RP0/CPU0:router(config-ospf)#router-id 2.2.2.2
```

```
RP/0/RP0/CPU0:router(config-ospf)#area 0
RP/0/RP0/CPU0:router(config-ospf-ar)#interface bundle-
RP/0/RP0/CPU0:router(config-ospf-ar-if)#cost-fallback 1000 threshold 300
```

Related Commands

| Command | Description |
|-----------------------------------------------|------------------------------------------------------------------------------|
| auto-cost (OSPF), on page 730 | Controls how the OSPF protocol calculates default metrics for the interface. |
| cost (OSPF), on page 740 | 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 803 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/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router ospf 1
RP/0/RP0/CPU0:router(config-ospf)# area 0
RP/0/RP0/CPU0:router(config-ospf-ar)# interface GigabitEthernet 0/1/0/1
RP/0/RP0/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 | <i>seconds</i> 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 | <p>If this command is not specified in interface configuration mode, then the interface adopts the dead interval parameter specified by the area.</p> <p>If this command is not specified in area configuration mode, then the interface adopts the dead interval parameter specified for the process.</p> <p>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.</p> | | | | |
| Command Modes | <p>Interface configuration</p> <p>Area configuration</p> <p>Router configuration</p> <p>Virtual-link configuration</p> <p>Multi-area configuration</p> <p>Sham-link configuration</p> | | | | |
| Command History | <table border="1"> <thead> <tr> <th>Release</th> <th>Modification</th> </tr> </thead> <tbody> <tr> <td>Release 5.0.0</td> <td>This command was introduced.</td> </tr> </tbody> </table> | Release | Modification | Release 5.0.0 | This command was introduced. |
| Release | Modification | | | | |
| Release 5.0.0 | This command was introduced. | | | | |
| Usage Guidelines | <p>To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.</p> <p>The dead interval value must be the same for all routers and access servers on a specific network.</p> <p>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.</p> | | | | |
| Task ID | <table border="1"> <thead> <tr> <th>Task ID</th> <th>Operations</th> </tr> </thead> <tbody> <tr> <td>ospf</td> <td>read, write</td> </tr> </tbody> </table> | Task ID | Operations | ospf | read, write |
| Task ID | Operations | | | | |
| ospf | read, write | | | | |

Examples

The following example shows how to set the OSPF dead interval to 40 seconds:

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router ospf 1
RP/0/RP0/CPU0:router(config-ospf)# area 0
RP/0/RP0/CPU0:router(config-ospf-ar)# interface GigabitEthernet 0/1/0/1
RP/0/RP0/CPU0:router(config-ospf-ar-if)# dead-interval 40
```

Related Commands

| Command | Description |
|----------------------------------------------------|-----------------------------------------------------------------------------------------------------|
| hello-interval (OSPF), on page 773 | 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 | <i>cost</i> Cost for the default summary route used for a stub or NSSA area. The acceptable value is a 24-bit number. | | | | |
|---------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------|--------------|---------------|------------------------------|
| Command Default | <i>cost</i> : 1 | | | | |
| Command Modes | Area configuration | | | | |
| Command History | <table border="1"> <thead> <tr> <th>Release</th> <th>Modification</th> </tr> </thead> <tbody> <tr> <td>Release 5.0.0</td> <td>This command was introduced.</td> </tr> </tbody> </table> | Release | Modification | Release 5.0.0 | This command was introduced. |
| Release | Modification | | | | |
| 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.

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/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router ospf 201
RP/0/RP0/CPU0:router(config-ospf)# area 10.15.0.0
RP/0/RP0/CPU0:router(config-ospf-ar)# stub
RP/0/RP0/CPU0:router(config-ospf-ar)# default-cost 20
RP/0/RP0/CPU0:router(config-ospf-ar)# interface GigabitEthernet 0/4/0/3
```

Related Commands

| Command | Description |
|------------------------------------------|---------------------------------|
| stub (OSPF), on page 891 | 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*] [**tag** *tag-value*]
no default-information originate

| Syntax Description | | |
|----------------------------------------|--|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| always | | (Optional) Always advertises the default route regardless of whether the routing table has a default route. |
| metric <i>metric-value</i> | | (Optional) Specifies the metric used for generating the default route. The default metric value is 1. Range is 1 to 16777214. |
| metric-type <i>type-value</i> | | (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 |
| tag <i>tag-value</i> | | (Optional) 32-bit dotted-decimal value attached to each external route. This is not used by the OSPF protocol itself. It may be used to communicate information between autonomous system boundary routers (ASBRs). If a tag is not specified, then the configured OSPF process number is used. |
| route-policy <i>policy-name</i> | | (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

tag-value: configured OSPF process number

Command Modes Router configuration

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | 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.

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 NCS 6000 Series Routers*.

For information about routing policies, see the *Routing Policy Commands* chapter in the *Routing Command Reference for Cisco NCS 6000 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/RP0/CPU0:router#configure
RP/0/RP0/CPU0:router(config)#router ospf 109
RP/0/RP0/CPU0:router(config-ospf)#redistribute igmp 108 metric 100
RP/0/RP0/CPU0:router(config-ospf)#default-information originate metric 100 metric-type 1
```

Related Commands

| Command | Description |
|--------------------------------------------------|-----------------------------------------------------------------------------|
| redistribute (OSPF), on page 832 | 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 | <i>value</i> 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 | | | | |
| Command History | <table border="1"> <thead> <tr> <th>Release</th> <th>Modification</th> </tr> </thead> <tbody> <tr> <td>Release 5.0.0</td> <td>This command was introduced.</td> </tr> </tbody> </table> | Release | Modification | Release 5.0.0 | This command was introduced. |
| Release | Modification | | | | |
| 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.

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/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router ospf 1
RP/0/RP0/CPU0:router(config-ospf)# default-metric 10
RP/0/RP0/CPU0:router(config-ospf)# redistribute isis IS-IS_isp
```

Related Commands

| Command | Description |
|--------------------------------------------------|-----------------------------------------------------------------------------|
| redistribute (OSPF), on page 832 | 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 | <p>disable (Optional) Disables the interface as an OSPF demand circuit.</p> <p>enable (Optional) Enables the interface as an OSPF demand circuit.</p> | | | | |
|---------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------|--------------|---------------|------------------------------|
| Command Default | <p>If this command is not specified in interface configuration mode, then the interface adopts the demand circuit parameter specified by the area.</p> <p>If this command is not specified in area configuration mode, then the interface adopts the demand circuit parameter specified for the process.</p> <p>If this command is not specified at any level, then the circuit is not a demand circuit.</p> | | | | |
| Command Modes | <p>Interface configuration</p> <p>Area configuration</p> <p>Router configuration</p> | | | | |
| Command History | <table border="1"> <thead> <tr> <th>Release</th> <th>Modification</th> </tr> </thead> <tbody> <tr> <td>Release 5.0.0</td> <td>This command was introduced.</td> </tr> </tbody> </table> | Release | Modification | Release 5.0.0 | This command was introduced. |
| Release | Modification | | | | |
| Release 5.0.0 | This command was introduced. | | | | |
| Usage Guidelines | <p>To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.</p> <p>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.</p> | | | | |
| Task ID | <table border="1"> <thead> <tr> <th>Task ID</th> <th>Operations</th> </tr> </thead> <tbody> <tr> <td>ospf</td> <td>read, write</td> </tr> </tbody> </table> | Task ID | Operations | ospf | read, write |
| Task ID | Operations | | | | |
| ospf | read, write | | | | |
| Examples | <p>The following example shows how to set the configuration for an OSPF demand circuit:</p> | | | | |

```
RP/0//CPU0:router# configure  
RP/0//CPU0:router(config)# router ospf 1  
RP/0//CPU0:router(config-ospf)# demand-circuit
```

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 | <i>weight</i> | Administrative distance. Range is 10 to 255. Used alone, the <i>weight</i> 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 75: Default Administrative Distances, on page 756 |
| | <i>ip-address</i> | (Optional) IP address in four-part, dotted-decimal notation. |
| | <i>wildcard-mask</i> | (Optional) Wildcard mask in four-part, dotted decimal format. A bit set to 1 in the <i>mask</i> argument instructs the software to ignore the corresponding bit in the address value. |
| | <i>access-list-name</i> | (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 75: Default Administrative Distances, on page 756 . | |
| Command Modes | Router configuration | |
| Command History | Release | Modification |
| | Release 5.0.0 | This command was introduced. |
| Usage Guidelines | <p>To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.</p> <p>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.</p> <p>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.</p> <p>The order in which you enter distance commands can affect the assigned administrative distances in unexpected ways (see the “Examples” section for further clarification).</p> | |

This table lists default administrative distances.

Table 75: 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 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/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router ospf 1
RP/0/RP0/CPU0:router(config-ospf)# distance 255
RP/0/RP0/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, on page 842 | 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. | | | | |
|---------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------|--------------|---------------|------------------------------|
| | <i>distance</i> Route administrative distance. | | | | |
| Command Default | <i>distance</i> : 110 | | | | |
| Command Modes | Router configuration | | | | |
| Command History | <table border="1"> <thead> <tr> <th>Release</th> <th>Modification</th> </tr> </thead> <tbody> <tr> <td>Release 5.0.0</td> <td>This command was introduced.</td> </tr> </tbody> </table> | Release | Modification | Release 5.0.0 | This command was introduced. |
| Release | Modification | | | | |
| 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.

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/RP0/CPU0:router# configure
```

```
RP/0/RP0/CPU0:router(config)# router ospf 1  
RP/0/RP0/CPU0:router(config-ospf)# redistribute ospf 2  
RP/0/RP0/CPU0:router(config-ospf)# distance ospf external 200  
RP/0/RP0/CPU0:router(config-ospf)# exit  
RP/0/RP0/CPU0:router(config)# router ospf 2  
RP/0/RP0/CPU0:router(config-ospf)# redistribute ospf 1  
RP/0/RP0/CPU0:router(config-ospf)# distance ospf external 200
```

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

| | |
|-------------------------|-------------------------------------------------------------------------------------------------------------------------------------|
| <i>access-list-name</i> | 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
 Multi-area configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| 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.

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/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# ipv4 access-list 3
RP/0/RP0/CPU0:router(config-ipv4-acl)# deny 172.17.10.0 0.0.0.255
RP/0/RP0/CPU0:router(config-ipv4-acl)# permit any any
!
RP/0/RP0/CPU0:router(config)# router ospf 1
RP/0/RP0/CPU0:router(config-ospf)# area 0
RP/0/RP0/CPU0:router(config-ospf-ar)# distribute-list 3 in
RP/0/RP0/CPU0:router(config-ospf-ar)# interface GigabitEthernet 0/1/0/3
```

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

| Command History | Release | Modification |
|-----------------|---------------|-----------------------------------------------------------------|
| | Release 5.0.0 | 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/RP0/CPU0:router (config)# router ospf 1
```

```
RP/0/RP0/CPU0:router(config-ospf)# area 0  
RP/0/RP0/CPU0:router(config-ospf-ar)# interface GigabitEthernet 0/3/0/0  
RP/0/RP0/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/RP0/CPU0:router#configure  
RP/0/RP0/CPU0:router(config)#router ospf 1  
RP/0/RP0/CPU0:router(config-ospf)#area 0  
RP/0/RP0/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/RP0/CPU0:router#configure  
RP/0/RP0/CPU0:router(config)#router ospf 1  
RP/0/RP0/CPU0:router(config-ospf)#area 0  
RP/0/RP0/CPU0:router(config-ospf-ar)#fast-reroute per-prefix  
RP/0/RP0/CPU0:router(config-ospf-ar)#interface GigabitEthernet 0/3/0/0  
RP/0/RP0/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 | <i>type</i> | Interface type. For more information, use the question mark (?) online help function. |
| | <i>interface-path-id</i> | 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

| | | |
|------------------------|----------------|------------------------------|
| Command History | Release | Modification |
| | 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 | 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/RP0/CPU0:router(config)# router ospf 1
RP/0/RP0/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 762 | 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 | <i>type</i> | Interface type. For more information, use the question mark (?) online help function. |
| | <i>interface-path-id</i> | 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

| | | |
|------------------------|----------------|------------------------------|
| Command History | Release | Modification |
| | 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.

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/RP0/CPU0:router#configure
RP/0/RP0/CPU0:router(config)#router ospf 100
RP/0/RP0/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 | <i>interface-name</i> 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 |
|----------------------|-----------------------------------------------------------------------|

| | | |
|------------------------|----------------|------------------------------|
| Command History | Release | Modification |
| | 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 |
| | ospf | read, write |

This example shows how to add an interface to LFA candidates:

```
RP/0/RP0/CPU0:router#configure
RP/0/RP0/CPU0:router(config)#router ospf 100
RP/0/RP0/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 <i>path-cost</i> | 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/RP0/CPU0:router(config)#router ospf 1
RP/0/RP0/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/RP0/CPU0:router(config)#router ospf 1
RP/0/RP0/CPU0:router(config-ospf)#fast-reroute per-prefix remote-lfa maximum-cost 2
```

Related Commands

| Command | Description |
|----------------------------------------------------|----------------------------------------------------------------|
| fast-reroute (OSPFv2), on page 762 | Enables IP fast reroute loop-free alternate (LFA) computation. |

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 | <p>enable Enables backup selection from candidate-list only.</p> <p>disable Disables backup selection from candidate-list only.</p> | | | | |
|---------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------|--------------|---------------|------------------------------|
| Command Default | Disabled. | | | | |
| Command Modes | Router OSPF configuration | | | | |
| Command History | <table border="1"> <thead> <tr> <th>Release</th> <th>Modification</th> </tr> </thead> <tbody> <tr> <td>Release 5.0.0</td> <td>This command was introduced.</td> </tr> </tbody> </table> | Release | Modification | Release 5.0.0 | This command was introduced. |
| Release | Modification | | | | |
| 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 | <table border="1"> <thead> <tr> <th>Task ID</th> <th>Operation</th> </tr> </thead> <tbody> <tr> <td>ospf</td> <td>read, write</td> </tr> </tbody> </table> | Task ID | Operation | ospf | read, write |
| 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/RP0/CPU0:router#configure
RP/0/RP0/CPU0:router(config)#router ospf 100
RP/0/RP0/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 | <p>enable (Optional) Turns on this functionality at a specific level.</p> <p>disable (Optional) Turns off this functionality at a specific level.</p> | | | | |
|---------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------|--------------|---------------|------------------------------|
| Command Default | <p>If this command is not specified in interface configuration mode, then the interface adopts the flood reduction parameter specified by the area.</p> <p>If this command is not specified in area configuration mode, then the interface adopts the flood reduction parameter specified for the process.</p> <p>If this command is not specified at any level, then flood reduction is disabled.</p> | | | | |
| Command Modes | <p>Interface configuration</p> <p>Area configuration</p> <p>Router configuration</p> | | | | |
| Command History | <table border="1"> <thead> <tr> <th>Release</th> <th>Modification</th> </tr> </thead> <tbody> <tr> <td>Release 5.0.0</td> <td>This command was introduced.</td> </tr> </tbody> </table> | Release | Modification | Release 5.0.0 | This command was introduced. |
| Release | Modification | | | | |
| Release 5.0.0 | This command was introduced. | | | | |
| Usage Guidelines | <p>To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.</p> <p>All routers supporting the OSPF demand circuit are compatible and can interact with routers supporting flooding reduction.</p> | | | | |
| Task ID | <table border="1"> <thead> <tr> <th>Task ID</th> <th>Operations</th> </tr> </thead> <tbody> <tr> <td>ospf</td> <td>read, write</td> </tr> </tbody> </table> | Task ID | Operations | ospf | read, write |
| Task ID | Operations | | | | |
| ospf | read, write | | | | |

Examples

The following example shows how to reduce the flooding of unnecessary LSAs for area 0:

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router (config)# router ospf 1
RP/0/RP0/CPU0:router (config-ospf)# area 0
RP/0/RP0/CPU0:router (config-ospf-ar)# interface GigabitEthernet 0/1/0/3
```

```
RP/0/RP0/CPU0:router (config-ospf-ar-if) # flood-reduction
```

Related Commands

| Command | Description |
|--------------------------------------------------|----------------------------------------------------------------------|
| show ospf interface, on page 853 | Displays OSPF-related interface information. |
| show ospf neighbor, on page 863 | 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 | <i>seconds</i> Interval (in seconds). The value must be the same for all nodes on a specific network. Range is 1 to 65535. | | | | |
|---------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------|--------------|---------------|------------------------------|
| Command Default | <p>If this command is not specified in interface configuration mode, then the interface adopts the hello interval parameter specified by the area.</p> <p>If this command is not specified in area configuration mode, then the interface adopts the hello interval parameter specified for the process.</p> <p>If this command is not specified at any level, then the hello interval is 10 seconds (broadcast) or 30 seconds (non-broadcast).</p> | | | | |
| Command Modes | <p>Interface configuration</p> <p>Area configuration</p> <p>Router configuration</p> <p>Virtual-link configuration</p> <p>Multi-area configuration</p> <p>Sham-link configuration</p> | | | | |
| Command History | <table border="1"> <thead> <tr> <th>Release</th> <th>Modification</th> </tr> </thead> <tbody> <tr> <td>Release 5.0.0</td> <td>This command was introduced.</td> </tr> </tbody> </table> | Release | Modification | Release 5.0.0 | This command was introduced. |
| Release | Modification | | | | |
| Release 5.0.0 | This command was introduced. | | | | |
| Usage Guidelines | <p>To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.</p> <p>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.</p> | | | | |
| Task ID | <table border="1"> <thead> <tr> <th>Task ID</th> <th>Operations</th> </tr> </thead> <tbody> <tr> <td>ospf</td> <td>read, write</td> </tr> </tbody> </table> | Task ID | Operations | ospf | read, write |
| Task ID | Operations | | | | |
| ospf | read, write | | | | |

Examples

The following example shows how to set the interval between hello packets to 15 seconds:

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router ospf 1
RP/0/RP0/CPU0:router(config-ospf)# area 0
RP/0/RP0/CPU0:router(config-ospf-ar)# interface GigabitEthernet 0/1/0/1
RP/0/RP0/CPU0:router(config-ospf-ar-if)# hello-interval 15
```

Related Commands

| Command | Description |
|---------------------------------------------------|------------------------------------------------------------------------------------------------------|
| dead-interval (OSPF), on page 745 | Sets the time period for which hello packets are suspended before neighbors declare the router down. |

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 **no interface** form of this command.

```
interface type interface-path-id
no interface type interface-path-id
```

| | | |
|---------------------------|--------------------------|------------------------------------------------------------------------------------------------------------|
| Syntax Description | <i>type</i> | Interface type. For more information, use the question mark (?) online help function. |
| | <i>interface-path-id</i> | 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, OSPF routing for interfaces is not enabled.

Command Modes Area configuration

| | | |
|------------------------|----------------|------------------------------|
| Command History | Release | Modification |
| | 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.

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/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router (config)# router ospf 109
RP/0/RP0/CPU0:router (config-ospf)# area 0
RP/0/RP0/CPU0:router (config-ospf-ar)# interface GigabitEthernet 4/0/0/3
!  
RP/0/RP0/CPU0:router (config-ospf)# area 2
```

```
RP/0/RP0/CPU0:router(config-ospf-ar)# interface GigabitEthernet 0/1/0/3
!
RP/0/RP0/CPU0:router(config-ospf)# area 3
RP/0/RP0/CPU0:router(config-ospf-ar)# interface GigabitEthernet 3/0/0/2
!
RP/0/RP0/CPU0:router(config-ospf)# area 10.9.50.0
RP/0/RP0/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 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.

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/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router ospf 109
RP/0/RP0/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 stub networks, 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 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.

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/RP0/CPU0:router(config)#router ospf 100
RP/0/RP0/CPU0:router(config-ospf)#loopback stub-network enable
```

| Related Commands | Command | Description |
|------------------|---------------------------------------------------|-----------------------------------------------------------------|
| | show ospf interface , on page 853 | Displays Open Shortest Path First (OSPF) interface information. |

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 | | |
|----------------------------------|--|--------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>max</i> | | Maximum number of nonself-generated LSAs the OSPF process can keep in the OSPF LSDB. |
| <i>threshold</i> | | (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 <i>value</i> | | (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 <i>value</i> | | (Optional) Specifies the number of times the OSPF process can consecutively be placed into the ignore state. The default is 5 times. |
| reset-time <i>value</i> | | (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

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | 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.

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

Related Commands

| Command | Description |
|----------------------------------------|---------------------------------------------------------------------------------------|
| show ospf, on page 846 | 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 <i>overriding metric</i> | (Optional) Overrides the external-lsa metric with the max-metric value. The <i>overriding metric</i> 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 <i>time</i> | (Optional) Sets the maximum metric temporarily after a process migration to originate router-LSAs with the max-metric value. The <i>time</i> range is 5 to 86400 seconds. |
| on-proc-restart <i>time</i> | (Optional) Sets the maximum metric temporarily after a process restart to originate router-LSAs with the max-metric value. The <i>time</i> range is 5 to 86400 seconds. |
| on-startup <i>time</i> | (Optional) Sets the maximum metric temporarily after a reboot to originate router-LSAs with the max-metric value. The <i>time</i> range is 5 to 86400 seconds. |
| on-switchover <i>time</i> | (Optional) Sets the maximum metric temporarily after a switchover to originate router-LSAs with the max-metric value. The <i>time</i> 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

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | 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.

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/RP0/CPU0:router# configure  
RP/0/RP0/CPU0:router(config)# router ospf 109  
RP/0/RP0/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 | <i>number-interfaces</i> Number of interfaces. Range is 1 to 4294967295. |
|---------------------------|--------------------------------------------------------------------------|

| | |
|------------------------|-------------------------------------------------------|
| Command Default | If the command is not specified, the default is 1024. |
|------------------------|-------------------------------------------------------|

| | |
|----------------------|----------------------|
| Command Modes | Router configuration |
|----------------------|----------------------|

| Command History | Release | Modification |
|------------------------|----------------|------------------------------|
| | Release 5.0.0 | This command was introduced. |

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

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/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router ospf 109
RP/0/RP0/CPU0:router(config-ospf)# maximum interfaces 1500
```

| Related Commands | Command | Description |
|-------------------------|--------------------------------------------------|--------------------------------------|
| | show ospf interface, on page 853 | 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 | <p><i>maximum-routes-number</i> Maximum number of parallel routes that OSPF can install in a routing table. Range is 1 to 64.</p> <p>Note The maximum number of paths that can be configured is 64.</p> |
|---------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

| | |
|------------------------|----------|
| Command Default | 64 paths |
|------------------------|----------|

| | |
|----------------------|----------------------|
| Command Modes | Router configuration |
|----------------------|----------------------|

| | | |
|------------------------|----------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Command History | Release | Modification |
| | Release 5.0.0 | This command was introduced. |
| | Release 6.1.2 | <p>This command was modified.</p> <p>The maximum number of paths that can be configured has increased from 32 to 64 paths.</p> <p>The command default has increased from 32 to 64 paths.</p> |

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

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 than 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/RP0/CPU0:router# configure  
RP/0/RP0/CPU0:router(config)# router ospf 109  
RP/0/RP0/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 | | |
|--------------------|------------------------|---------------------------------------------------------------------------------------------------------|
| | <i>maximum</i> | Number of routes. Range is 1 to 4294967295. |
| | <i>threshold-value</i> | (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

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | 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.

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 is less than the existing number of routes, existing routes remain configured, but no new routes are redistributed.

| 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/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router ospf 109
RP/0/RP0/CPU0:router(config-ospf)# maximum redistributed-prefixes 15000
```

Related Commands

| Command | Description |
|-----------------------------------------------|-----------------------------------|
| show ospf routes, on page 874 | 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

| | |
|------------------|----------------------------------------------------------------|
| <i>key-id</i> | Key number. Range is 1 to 255. |
| md5 | Enables OSPF MD5 authentication. |
| <i>key</i> | 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
 Multi-area configuration
 Sham-link configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| 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.

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/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router ospf 109
RP/0/RP0/CPU0:router(config-ospf)# area 0

RP/0/RP0/CPU0:router(config-ospf-ar)# interface GigabitEthernet 0/1/0/1
RP/0/RP0/CPU0:router(config-ospf-ar-if)# message-digest-key 19 md5 8ry4222
```

Related Commands

| Command | Description |
|----------------------------------------------------|---------------------------------------------------------------------------------------|
| area (OSPF), on page 723 | Configures an OSPF area. |
| authentication (OSPF), on page 725 | Enables plain text, MD5 authentication, or null authentication for an OSPF interface. |
| default-cost (OSPF), on page 747 | Enables authentication for an OSPF area. |

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 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 | Operations |
|---------|---------|----------------|
| | ospf | read, write |

Examples

The following example shows how to enable LDP-IGP interface automatic configuration:

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router ospf 01
RP/0/RP0/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 | <table border="1"> <thead> <tr> <th>Release</th> <th>Modification</th> </tr> </thead> <tbody> <tr> <td>Release 5.0.0</td> <td>This command was introduced.</td> </tr> </tbody> </table> | Release | Modification | Release 5.0.0 | This command was introduced. |
| Release | Modification | | | | |
| 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 | <table border="1"> <thead> <tr> <th>Task ID</th> <th>Operations</th> </tr> </thead> <tbody> <tr> <td>ospf</td> <td>read, write</td> </tr> </tbody> </table> | Task ID | Operations | ospf | read, write |
| Task ID | Operations | | | | |
| ospf | read, write | | | | |

Examples

The following example shows how to enable LDP-IGP synchronization:

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router ospf 01
RP/0/RP0/CPU0:router(config-ospf)# mpls ldp sync
```

mpls ldp sync-igp-shortcuts

To configure and enable Label Distribution Protocol- Interior Gateway Protocol (LDP-IGP) synchronization on Multiprotocol Label Switching (MPLS) shortcuts (auto-route announced tunnels) only, use the **mpls ldp sync-igp-shortcuts** command in the appropriate OSPF configuration mode. To disable the use of LDP-IGP synchronization on MPLS shortcuts, use the **no** form of this command.

mpls ldp sync-igp-shortcuts [**disable**]
no mpls ldp sync-igp-shortcuts

| | |
|---------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax Description | disable (Optional) Disables MPLS LDP synchronization with IGP shortcuts 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 on MPLS tunnels. |
|------------------------|------------------------------------------------------|

| | |
|----------------------|--------------------------------------------------------------------------------------|
| Command Modes | OSPF interface configuration OSPF router configuration OSPF area configuration |
|----------------------|--------------------------------------------------------------------------------------|

| Command History | Release | Modification |
|------------------------|----------------|------------------------------|
| | 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. |
|-------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

When **mpls ldp sync-igp-shortcuts** is configured at a higher mode (area or instance level), you can selectively disable LDP-IGP synchronization on specific interfaces or areas by using the **disable** form of the command.

| Task ID | Task ID | Operations |
|----------------|----------------|-------------------|
| | ospf | read, write |

| | |
|-----------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Examples | The following example shows how to enable LDP-IGP synchronization shortcuts under OSPF interface configuration, using the mpls ldp sync-igp-shortcuts command on a tunnel-te interface. However, the command applies to all interfaces under the applicable configuration mode. |
|-----------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

```
RP/0/RP0/CPU0:router (config) # router ospf 100
RP/0/RP0/CPU0:router (config-ospf) # area 0
RP/0/RP0/CPU0:router (config-ospf-ar) # interface tunnel-te 1
RP/0/RP0/CPU0:router (config-ospf-ar-if) # mpls ldp sync-igp-shortcuts
```

When **mpls ldp sync-igp-shortcuts** is configured at a higher mode (area or instance level), you can selectively disable LDP-IGP synchronization on specific interfaces or areas by using the **disable** form of the command:

The following example shows how to enable the use of LDP-IGP synchronization across all interfaces in all areas except area 10:

```
RP/0/RP0/CPU0:router(config)# router ospf foo
RP/0/RP0/CPU0:router(config-ospf)# mpls ldp sync-igp-shortcuts
RP/0/RP0/CPU0:router(config-ospf)# area 10
RP/0/RP0/CPU0:router(config-ospf-ar)# mpls ldp sync-igp-shortcuts disable
```

Related Commands

| Command | Description |
|---------------------------------------------------|--------------------------------------------------------------------------------------------|
| mpls ldp sync (OSPF), on page 793 | Enables Label Distribution Protocol (LDP)-Interior Gateway Protocol (IGP) synchronization. |

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

Usage Guidelines To use this command, you must be in a user group associated with a task 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/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router ospf 01
RP/0/RP0/CPU0:router(config-ospf)# mpls traffic-eng ldp-sync-update
```


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

Multi-area configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| 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.

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/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router ospf 109
RP/0/RP0/CPU0:router(config-ospf)# area 0
RP/0/RP0/CPU0:router(config-ospf-ar)# interface GigabitEthernet 0/1/0/3
RP/0/RP0/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 | <i>type</i> | Interface type. For more information, use the question mark (?) online help function. |
| | <i>interface-path-id</i> | 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 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.

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 OSPF speakers are attached. In the case of native broadcast networks, the interface must be configured as an OSPF 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/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router ospf 109
RP/0/RP0/CPU0:router(config-ospf)# area 0
RP/0/RP0/CPU0:router(config-ospf-ar)# interface GigabitEthernet 0/1/0/3
RP/0/RP0/CPU0:router(config-ospf-ar-if)# area 1
RP/0/RP0/CPU0:router(config-ospf-ar)# multi-area-interface GigabitEthernet 0/1/0/3
RP/0/RP0/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/RP0/CPU0:router(config-ospf-ar-mif)#

```

Related Commands

| Command | Description |
|--------------------------------------------------|--------------------------------------|
| show ospf interface, on page 853 | 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 | |
|-------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>ip-address</i> | Interface IP address of the neighbor. |
| cost <i>number</i> | (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 <i>number</i> | (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 <i>seconds</i> | (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 <i>number</i> : 0 poll-interval <i>seconds</i> : 120 seconds (2 minutes) |
|-----------------|--------------------------------------------------------------------------------------------------------------------------------------|
|-----------------|--------------------------------------------------------------------------------------------------------------------------------------|

| Command Modes | Interface configuration |
|---------------|-------------------------|
|---------------|-------------------------|

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | 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.

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/RP0/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/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# interface GigabitEthernet1/0/0/3
RP/0/RP0/CPU0:router(config-if)# ip address 172.16.3.10 255.255.255.0

RP/0/RP0/CPU0:router(config)# router ospf 1
RP/0/RP0/CPU0:router(config-ospf)# area 0
RP/0/RP0/CPU0:router(config-ospf-ar)# interface GigabitEthernet1/0/0/3
RP/0/RP0/CPU0:router(config-ospf-ar-if)# network nonbroadcast
RP/0/RP0/CPU0:router(config-ospf-ar-if)# neighbor 172.16.3.4 priority 1 poll-interval 180

RP/0/RP0/CPU0:router(config-ospf-ar-if)# neighbor 172.16.3.5 cost 10 priority 1
poll-interval 180
RP/0/RP0/CPU0:router(config-ospf-ar-if)# neighbor 172.16.3.6 cost 15 priority 1
poll-interval 180
RP/0/RP0/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 803 | Filters all outgoing LSAs to an OSPF neighbor. |
| network (OSPF), on page 804 | Configures the OSPF network type to a type other than the default for a given medium. |
| priority (OSPF), on page 818 | 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 | <i>ip-address</i> 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 | <table border="1"> <thead> <tr> <th>Release</th> <th>Modification</th> </tr> </thead> <tbody> <tr> <td>Release 5.0.0</td> <td>This command was introduced.</td> </tr> </tbody> </table> | Release | Modification | Release 5.0.0 | This command was introduced. |
| Release | Modification | | | | |
| 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.

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/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router (config)# router ospf 1
RP/0/RP0/CPU0:router (config-ospf)# area 0
RP/0/RP0/CPU0:router (config-ospf-ar)# interface GigabitEthernet1/0/0/3
RP/0/RP0/CPU0:router (config-ospf-ar-if)# neighbor 10.2.3.4 database-filter all out
```

| Related Commands | Command | Description |
|------------------|----------------------------------------------|-------------------------------------------------------------------|
| | neighbor (OSPF), on page 801 | 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. POS interfaces default to point-to-point and GigabitEthernet and TenGigEthernet interfaces default to broadcast.

Command Modes

Interface configuration
 Area configuration
 Router configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| 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.

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//CPU0:router# configure
RP/0//CPU0:router(config)# router ospf 1
RP/0//CPU0:router(config-ospf)# area 0
RP/0//CPU0:router(config-ospf-ar)# interface GigabitEthernet 0/1/0/3
RP/0//CPU0:router(config-ospf-ar-if)# network non-broadcast
RP/0//CPU0:router(config-ospf-ar-if)# neighbor 172.16.3.4 priority 1 poll-interval 180
```

Related Commands

| Command | Description |
|----------------------------------------------|-------------------------------------------------------------------|
| neighbor (OSPF), on page 801 | 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 nsf** 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 | |
| Command History | Release | Modification |
| | 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.

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/RP0/CPU0:router# configure  
RP/0/RP0/CPU0:router(config)# router ospf 1  
RP/0/RP0/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 | <i>seconds</i> Length of time (in seconds) allowed for NSF external route queries. Range is 1 to 3600 seconds. | |
| Command Default | <i>seconds</i> : 300 | |
| Command Modes | Router configuration | |
| Command History | Release | Modification |
| | 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 | 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/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router ospf 1
RP/0/RP0/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 | <i>seconds</i> Length of time (in seconds) between consecutive restart attempts. Range is 90 to 3600 seconds. | | | | |
|---------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------|--------------|---------------|------------------------------|
| Command Default | <i>seconds</i> : 90 | | | | |
| Command Modes | Router configuration | | | | |
| Command History | <table border="1"> <thead> <tr> <th>Release</th> <th>Modification</th> </tr> </thead> <tbody> <tr> <td>Release 5.0.0</td> <td>This command was introduced.</td> </tr> </tbody> </table> | Release | Modification | Release 5.0.0 | This command was introduced. |
| Release | Modification | | | | |
| 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.

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/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router ospf 1
RP/0/RP0/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 | <i>seconds</i> The length of time (in seconds) that routes are held in the RIB. Range is 90 to 3600 seconds. |
|---------------------------|--------------------------------------------------------------------------------------------------------------|

| | |
|------------------------|---------------------|
| Command Default | <i>seconds</i> : 95 |
|------------------------|---------------------|

| | |
|----------------------|----------------------|
| Command Modes | Router configuration |
|----------------------|----------------------|

| Command History | Release | Modification |
|------------------------|----------------|------------------------------|
| | 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. |
|-------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

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/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router ospf 1
RP/0/RP0/CPU0:router(config-ospf)# nsf lifetime 120
```

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 <i>metric-value</i> | (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 <i>type-value</i> | (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 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.

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/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router ospf 1
RP/0/RP0/CPU0:router(config-ospf)# area 1
RP/0/RP0/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 XR Config 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 XR Config

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | 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.

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/RP0/CPU0:router# configure
RP/0/RP0/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 | <i>bytes</i> 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 Multi-area configuration | |
| Command History | Release | Modification |
| | Release 5.0.0 | This command was introduced. |
| Usage Guidelines | <p>To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.</p> <p>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.</p> <p>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.</p> | |
| Task ID | Task ID | Operations |
| | ospf | read, write |
| Examples | The following example shows how to configure the packet size on an interface: | |

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

```
RP/0/RP0/CPU0:router(config)# router ospf 1  
RP/0/RP0/CPU0:router(config-ospf)# area 0  
RP/0/RP0/CPU0:router(config-ospf-ar)# interface GigabitEthernet 1/0/0/2  
RP/0/RP0/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

Command History

Release

Release 5.0.0

Modification

This command was introduced.

Usage Guidelines

To use this command, you must be in a user group associated with a task 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 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/RP0/CPU0:router# configure
```

```
RP/0/RP0/CPU0:router(config)# router ospf 1  
RP/0/RP0/CPU0:router(config-ospf)# area 0  
RP/0/RP0/CPU0:router(config-ospf-ar)# interface GigabitEthernet1/0/0/2  
RP/0/RP0/CPU0:router(config-ospf-ar-if)# passive  
RP/0/RP0/CPU0:router(config-ospf-ar-if)# exit  
RP/0/RP0/CPU0:router(config-ospf-ar)# interface GigabitEthernet1/0/0/3  
RP/0/RP0/CPU0:router(config-ospf-ar-if)# end
```

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 | <i>value</i> 8-bit unsigned integer indicating the router priority value. Range is 0 to 255. |
|---------------------------|----------------------------------------------------------------------------------------------|

| | |
|------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Command Default | <p>If this command is not specified in interface configuration mode, then the interface adopts the priority parameter specified by the area.</p> <p>If this command is not specified in area configuration mode, then the interface adopts the priority parameter specified for the process.</p> <p>If this command is not specified at any level, then the default priority is 1.</p> |
|------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

| | |
|----------------------|--------------------------------------------------------------------------------------|
| Command Modes | <p>Interface configuration</p> <p>Area configuration</p> <p>Router configuration</p> |
|----------------------|--------------------------------------------------------------------------------------|

| | | |
|------------------------|----------------|------------------------------|
| Command History | Release | Modification |
| | 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. |
|-------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

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/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# interface GigabitEthernet 0/1/0/1
RP/0/RP0/CPU0:router(config-if)# ipv4 address 10.0.0.2 255.255.255.0
RP/0/RP0/CPU0:router(config-if)# exit
RP/0/RP0/CPU0:router(config)# router ospf 1
RP/0/RP0/CPU0:router(config-ospf)# area 0
RP/0/RP0/CPU0:router(config-ospf-ar)# interface GigabitEthernet 0/1/0/1
RP/0/RP0/CPU0:router(config-ospf-ar-if)# network non-broadcast
RP/0/RP0/CPU0:router(config-ospf-ar-if)# priority 4
RP/0/RP0/CPU0:router(config-ospf-ar-if)# neighbor 10.0.0.1 priority 6

```

```

RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# interface GigabitEthernet POS 0/2/0/1
RP/0/RP0/CPU0:router(config-if)# ipv4 address 10.0.0.1 255.255.255.0
RP/0/RP0/CPU0:router(config-if)# exit
RP/0/RP0/CPU0:router(config)# router ospf 1
RP/0/RP0/CPU0:router(config-ospf)# area 0
RP/0/RP0/CPU0:router(config-ospf-ar)# interface GigabitEthernet 0/2/0/1
RP/0/RP0/CPU0:router(config-ospf-ar-if)# network non-broadcast
RP/0/RP0/CPU0:router(config-ospf-ar-if)# priority 6
RP/0/RP0/CPU0:router(config-ospf-ar-if)# neighbor 10.0.0.2 priority 4

```

Related Commands

| Command | Description |
|----------------------------------------------|---------------------------------------------------------------------------------------|
| neighbor (OSPF), on page 801 | Configures OSPF routers interconnecting to nonbroadcast networks. |
| network (OSPF), on page 804 | 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 reenble 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 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.

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/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router (config)# router ospfv3 1
RP/0/RP0/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 | <i>count</i> 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 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 | Operations |
| | ospf | read, write |

| | |
|-----------------|------------------------------------------------------------------------------------|
| Examples | The following example shows how to limit the number of LSAs flushed per run to 30: |
|-----------------|------------------------------------------------------------------------------------|

```
RP/0/RP0/CPU0:router(config-ospf)# queue dispatch flush-lsa 30
```

Use the [show ospf message-queue, on page 860](#) command to see the queue dispatch values, peak lengths, and limits.

| | | |
|-------------------------|--------------------------------------------------------------|---------------------------------------------------------------------------------------------|
| Related Commands | Command | Description |
| | queue dispatch incoming, on page 823 | Limits the number of continuous incoming events processed. |
| | queue dispatch rate-limited-lsa, on page 825 | Sets the maximum number of rate-limited link-state advertisements (LSAs) processed per run. |

| Command | Description |
|-----------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------|
| queue dispatch spf-lsa-limit, on page 827 | 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 828 | Sets the high watermark for incoming priority events. |
| show ospf message-queue, on page 860 | Displays the information about the queue dispatch values, peak lengths, and limits. |

queue dispatch incoming

To limit the number of incoming packets (LSAUpdates, LSAsacks, 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 | <i>count</i> 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 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 | Operations |
| | ospf | read, write |

| | |
|-----------------|----------------------------------------------------------------------------------------|
| Examples | The following example shows how limit the number of incoming packets processed to 500: |
|-----------------|----------------------------------------------------------------------------------------|

```
RP/0/RP0/CPU0:router (config-ospf) # queue dispatch incoming 500
```

Use the [show ospf message-queue, on page 860](#) command to see the queue dispatch values, peak lengths, and limits.

| Related Commands | Command | Description |
|------------------|--------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------|
| | queue dispatch rate-limited-lsa, on page 825 | Sets the maximum number of rate-limited link-state advertisements (LSAs) processed per run. |
| | queue dispatch spf-lsa-limit, on page 827 | 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 828 | Sets the high watermark for incoming priority events. |
| show ospf message-queue, on page 860 | 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 | <i>count</i> 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 | <table border="1"> <thead> <tr> <th>Release</th> <th>Modification</th> </tr> </thead> <tbody> <tr> <td>Release 5.0.0</td> <td>This command was introduced.</td> </tr> </tbody> </table> | Release | Modification | Release 5.0.0 | This command was introduced. | | | | |
| Release | Modification | | | | | | | | |
| 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 | <table border="1"> <thead> <tr> <th>Task ID</th> <th>Operations</th> </tr> </thead> <tbody> <tr> <td>ospf</td> <td>read, write</td> </tr> </tbody> </table> | Task ID | Operations | ospf | read, write | | | | |
| Task ID | Operations | | | | | | | | |
| ospf | read, write | | | | | | | | |
| Examples | <p>The following example shows how to set the maximum number of rate-limited LSA (re-)originations processed per run to 300:</p> <pre>RP/0/RP0/CPU0:router (config-ospf) # queue dispatch rate-limited-lsa 300</pre> | | | | | | | | |
| Related Commands | <table border="1"> <thead> <tr> <th>Command</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>queue dispatch incoming, on page 823</td> <td>Limits the number of continuous incoming events processed.</td> </tr> <tr> <td>queue dispatch spf-lsa-limit, on page 827</td> <td>Limits the number of summary or external Type 3 to Type 7 link-state advertisements (LSAs) processed per shortest path first (SPF) run.</td> </tr> <tr> <td>queue limit, on page 828</td> <td>Sets the high watermark for incoming priority events.</td> </tr> </tbody> </table> | Command | Description | queue dispatch incoming, on page 823 | Limits the number of continuous incoming events processed. | queue dispatch spf-lsa-limit, on page 827 | 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 828 | Sets the high watermark for incoming priority events. |
| Command | Description | | | | | | | | |
| queue dispatch incoming, on page 823 | Limits the number of continuous incoming events processed. | | | | | | | | |
| queue dispatch spf-lsa-limit, on page 827 | 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 828 | Sets the high watermark for incoming priority events. | | | | | | | | |

| Command | Description |
|------------------------------------------------------|-------------------------------------------------------------------------------------|
| show ospf message-queue, on page 860 | 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 | <i>count</i> 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 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 | 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/RP0/CPU0:router(config-ospf)# queue dispatch spf-lsa-limit 100
```

| | | |
|-------------------------|--------------------------------------------------------------|--------------------------------------------------------------------------------------------|
| Related Commands | Command | Description |
| | queue dispatch incoming, on page 823 | Limits the number of continuous incoming events processed. |
| | queue dispatch rate-limited-lsa, on page 825 | Sets the maximum number of rate-limited link-state advertisements (LSAs) processed per run |
| | queue limit, on page 828 | Sets the high watermark for incoming priority events. |
| | show ospf message-queue, on page 860 | 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). |
| <i>count</i> | 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 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.

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/RP0/CPU0:router(config-ospf)# queue limit high 11000
RP/0/RP0/CPU0:router(config-ospf)# queue limit medium 10000
RP/0/RP0/CPU0:router(config-ospf)# queue limit low 9000
```


Related Commands

| Command | Description |
|--------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------|
| queue dispatch incoming, on page 823 | Limits the number of continuous incoming events processed. |
| queue dispatch rate-limited-lsa, on page 825 | Sets the maximum number of rate-limited link-state advertisements (LSAs) processed per run. |
| queue dispatch spf-lsa-limit, on page 827 | 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 860 | 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

| | |
|----------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>ip-address</i> | IP address in four-part, dotted-decimal notation. |
| <i>mask</i> | 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 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.

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/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router (config)# router ospf 201
RP/0/RP0/CPU0:router (config-ospf)# area 0
RP/0/RP0/CPU0:router (config-ospf-ar-if)# interface GigabitEthernet 0/3/0/2
!
RP/0/RP0/CPU0:router (config-ospf)# area 36.0.0.0
RP/0/RP0/CPU0:router (config-ospf-ar)# range 10.31.0.0 255.255.0.0
RP/0/RP0/CPU0:router (config-ospf-ar)# interface GigabitEthernet 0/1/0/0
RP/0/RP0/CPU0:router (config-ospf-ar-if)# interface GigabitEthernet0/1/0/0
RP/0/RP0/CPU0:router (config-ospf-ar-if)# interface GigabitEthernet0/1/0/1
RP/0/RP0/CPU0:router (config-ospf-ar-if)# interface GigabitEthernet0/1/0/2
RP/0/RP0/CPU0:router (config-ospf-ar-if)# interface GigabitEthernet0/1/0/3
RP/0/RP0/CPU0:router (config-ospf-ar-if)# interface GigabitEthernet0/2/0/0
RP/0/RP0/CPU0:router (config-ospf-ar-if)# interface GigabitEthernet0/2/0/1
RP/0/RP0/CPU0:router (config-ospf-ar-if)# interface GigabitEthernet0/2/0/2
RP/0/RP0/CPU0:router (config-ospf-ar-if)# interface GigabitEthernet0/2/0/3
RP/0/RP0/CPU0:router (config-ospf-ar-if)# end
```

Related Commands

| Command | Description |
|----------------------------------------------------|------------------------------------------------------------------------------------------------------------------|
| summary-prefix (OSPF), on page 893 | 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*]

| Syntax Description | |
|--------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| bgp | Distributes routes from the BGP protocol. |
| <i>process-id</i> | For the bgp keyword, an autonomous system number has the following ranges: <ul style="list-style-type: none"> • 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. 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 <i>metric-value</i> | (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: <ul style="list-style-type: none"> • 1 —Type 1 external route • 2 —Type 2 external route |
| tag <i>tag-value</i> | (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 <i>policy-name</i> | (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. |
| <i>instance-name</i> | 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. |
| level-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] } | <p>(Optional) Specifies the criteria by which OSPF routes are redistributed into other routing domains. It can be one or more of the following:</p> <ul style="list-style-type: none"> • 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. <p>For the external and nssa-external options, if a type is not specified, then both Type 1 and Type 2 are assumed.</p> <p>If no match is specified, the default is no filtering.</p> |
| rip | Distributes routes from the RIP protocol. |
| static | Distributes IP static routes. |
| dagr | Distributes routes from the directed-attached gateway redundancy (DAGR). |

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

Command History

| Release | Modification |
|---------------|------------------------------|
| 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.



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 module of *Routing Command Reference for Cisco NCS 6000 Series Routers*.

Whenever you use the **redistribute** or [default-information originate \(OSPF\)](#), on page 749 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/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router ospf 110
RP/0/RP0/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/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router ospf 109
RP/0/RP0/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/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# interface GigabitEthernet 0/1/0/1
RP/0/RP0/CPU0:router(config-if)# ip address 10.0.0.0 255.0.0.0
!
RP/0/RP0/CPU0:router(config)# interface GigabitEthernet 0/2/0/2
RP/0/RP0/CPU0:router(config)# ip address 10.99.0.0 255.0.0.0
!
RP/0/RP0/CPU0:router(config)# router ospf 1
RP/0/RP0/CPU0:router(config-ospf)# redistribute ospf 2
RP/0/RP0/CPU0:router(config-ospf)# area 0
RP/0/RP0/CPU0:router(config-ospf-ar)# interface GigabitEthernet 0/2/0/2
```

redistribute (OSPF)

```
!  
RP/0/RP0/CPU0:router(config)# router ospf 2  
RP/0/RP0/CPU0:router(config-ospf)# area 0  
RP/0/RP0/CPU0:router(config-ospf-ar)# interface GigabitEthernet 0/1/0/1
```

Related Commands

| Command | Description |
|-------------------------------------------------------------------|-----------------------------------------------------------------|
| default-information originate (OSPF), on page 749 | 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 | <i>seconds</i> 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 | <p>If this command is not specified in interface configuration mode, then the interface adopts the retransmit interval parameter specified by the area.</p> <p>If this command is not specified in area configuration mode, then the interface adopts the retransmit interval parameter specified for the process.</p> <p>If this command is not specified at any level, then the default retransmit interval is 5 seconds.</p> | | | | |
| Command Modes | <p>Interface configuration</p> <p>Area configuration</p> <p>Router configuration</p> <p>Virtual-link configuration</p> <p>Multi-area configuration</p> <p>Sham-link configuration</p> | | | | |
| Command History | <table border="1"> <thead> <tr> <th>Release</th> <th>Modification</th> </tr> </thead> <tbody> <tr> <td>Release 5.0.0</td> <td>This command was introduced.</td> </tr> </tbody> </table> | Release | Modification | Release 5.0.0 | This command was introduced. |
| Release | Modification | | | | |
| Release 5.0.0 | This command was introduced. | | | | |
| Usage Guidelines | <p>To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.</p> <p>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.</p> <p>The setting of this parameter should be conservative, or needless retransmission results. The value should be larger for serial lines and virtual links.</p> | | | | |
| Task ID | <table border="1"> <thead> <tr> <th>Task ID</th> <th>Operations</th> </tr> </thead> <tbody> <tr> <td>ospf</td> <td>read, write</td> </tr> </tbody> </table> | Task ID | Operations | ospf | read, write |
| 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/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router ospf 201
RP/0/RP0/CPU0:router(config-ospf)# area 0
RP/0/RP0/CPU0:router(config-ospf-ar)# interface GigabitEthernet 0/2/0/1
RP/0/RP0/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 | <table border="1"> <tr> <td><i>route-policy-name</i></td> <td>Name of route policy.</td> </tr> <tr> <td>in</td> <td>Applies policy to inbound routes.</td> </tr> <tr> <td>out</td> <td>Applies policy to outbound routes.</td> </tr> </table> | <i>route-policy-name</i> | Name of route policy. | in | Applies policy to inbound routes. | out | Applies policy to outbound routes. |
|---------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------|-----------------------|---------------|-----------------------------------|------------|------------------------------------|
| <i>route-policy-name</i> | 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 | <table border="1"> <thead> <tr> <th>Release</th> <th>Modification</th> </tr> </thead> <tbody> <tr> <td>Release 5.0.0</td> <td>This command was introduced.</td> </tr> </tbody> </table> | Release | Modification | Release 5.0.0 | This command was introduced. | | |
| Release | Modification | | | | | | |
| Release 5.0.0 | This command was introduced. | | | | | | |
| Usage Guidelines | <p>To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.</p> <p>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.</p> | | | | | | |
| Task ID | <table border="1"> <thead> <tr> <th>Task ID</th> <th>Operations</th> </tr> </thead> <tbody> <tr> <td>ospf</td> <td>read, write</td> </tr> </tbody> </table> | Task ID | Operations | ospf | read, write | | |
| 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/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router ospf 109
RP/0/RP0/CPU0:router(config-ospf)# area 0
RP/0/RP0/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 | <i>router-id</i> 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 | | | | |
| Command History | <table border="1"> <thead> <tr> <th>Release</th> <th>Modification</th> </tr> </thead> <tbody> <tr> <td>Release 5.0.0</td> <td>This command was introduced.</td> </tr> </tbody> </table> | Release | Modification | Release 5.0.0 | This command was introduced. |
| Release | Modification | | | | |
| 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.

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/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router ospf 109
RP/0/RP0/CPU0:router(config-ospf)# router-id 172.20.10.10
```

Related Commands

| Command | Description |
|-------------------------------------------------|-------------------------------------------------------------------|
| clear ospf process, on page 733 | Resets an OSPF router process without stopping and restarting it. |
| ipv4 address | Sets a primary IPv4 address for an interface. |

router ospf

To configure an Open Shortest Path First (OSPF) routing process, use the **router ospf** command in XR Config 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

XR Config

Command History

| Release | Modification |
|---------------|------------------------------|
| 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.

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/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router ospf 109
```

Related Commands

| Command | Description |
|------------------------------------------|--------------------------|
| area (OSPF), on page 723 | Configures an OSPF area. |

| Command | Description |
|----------------------------------------------------|---------------------------------------------------------------------------------------------------|
| default-metric (OSPF), on page 751 | Sets default metric values for routes redistributed from another protocol into the OSPF protocol. |
| interface (OSPF), on page 775 | Defines the interfaces on which the OSPF protocol runs. |
| router-id (OSPF), on page 840 | 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 <i>hops-number</i> IP hops. Maximum number of hops allowed. Range is 1 to 254 hops. |
|---------------------------|-------------------------------------------------------------------------------------------------|

| | |
|------------------------|------------------------|
| Command Default | <i>hops-number</i> : 1 |
|------------------------|------------------------|

| | |
|----------------------|-----------------------------------------------------------------------|
| Command Modes | Router configuration Area configuration Interface configuration |
|----------------------|-----------------------------------------------------------------------|

| | | |
|------------------------|----------------|------------------------------|
| Command History | Release | Modification |
| | 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.

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 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/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router ospf 1
RP/0/RP0/CPU0:router(config-ospf)# area 0
RP/0/RP0/CPU0:router(config-ospf-ar)# interface GigabitEthernet0/6/0/3
RP/0/RP0/CPU0:router(config-ospf-ar-if)# security ttl 2
```

show ospf

To display general information about Open Shortest Path First (OSPF) routing processes, use the **show ospf** command in XR EXEC mode.

show ospf [*process-name*] [**summary**]

| | |
|---------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax Description | <i>process-name</i> (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. |
| | summary (Optional) Displays OSPF summary information. |

Command Default IPv4 and unicast address prefixes

Command Modes XR EXEC
OSPFv3

| Command History | Release | Modification |
|------------------------|----------------|------------------------------|
| | 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.

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/RP0/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 msec
  Minimum hold time between two consecutive SPF's 10000 msec
  Maximum wait time between two consecutive SPF's 10000 msec
  Initial LSA throttle delay 500 msec
  Minimum hold time for LSA throttle 5000 msec
  Maximum wait time for LSA throttle 5000 msec
  Minimum LSA interval 5000 msec. Minimum LSA arrival 1 sec
  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 76: 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. |
| 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. |

show ospf

| Field | Description |
|---------------|---------------------|
| Area BACKBONE | Backbone is area 0. |

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 XR EXEC mode.

```
show ospf [process-name] border-routers [router-id]
```

| Syntax Description | |
|---------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>process-name</i> | (Optional) OSPF process name. If this argument is included, only information for the specified routing process is included. |
| <i>router-id</i> | (Optional) Router ID associated with the border router. The value of the <i>router-id</i> 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 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.

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/RP0/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 77: 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 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 XR EXEC mode.

```
show ospf [process-name] [area-id] flood-list [type interface-path-id]
```

| Syntax Description | |
|--------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>process-name</i> | (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. |
| <i>area-id</i> | (Optional) Area number used to define the particular area. |
| <i>type</i> | Interface type. For more information, use the question mark (?) online help function. |
| <i>interface-path-id</i> | Physical interface or virtual interface. |
| | <p>Note Use the show interfaces command to see a list of all interfaces currently configured on the router.</p> <p>For more information about the syntax for the router, use the question mark (?) online help function.</p> |

| | |
|------------------------|----------------|
| Command Default | All interfaces |
|------------------------|----------------|

| | |
|----------------------|---------|
| Command Modes | XR EXEC |
|----------------------|---------|

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | 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.

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/RP0/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 78: 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 XR EXEC mode.

```
show ospf [process-name] [area-id] interface [brief] [type interface-path-id]
```

| Syntax Description | |
|--------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>process-name</i> | (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. |
| <i>area-id</i> | (Optional) Area number used to define the particular area. |
| brief | (Optional) Displays brief interface information. |
| <i>type</i> | Interface type. For more information, use the question mark (?) online help function. |
| <i>interface-path-id</i> | 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 | XR EXEC |
|----------------------|---------|

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | 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 | 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/RP0/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
```

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 XR EXEC mode.

```
show ospf [process-name] [area-id] [type interface-path-id] mpls traffic-eng {link | fragment}
```

| Syntax Description | |
|--------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>process-name</i> | (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. |
| <i>area-id</i> | (Optional) Area number used to define the particular area. |
| <i>type</i> | Interface type. For more information, use the question mark (?) online help function. |
| <i>interface-path-id</i> | 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 XR EXEC

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | 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 | 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/RP0/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 79: show ospf mpls traffic-eng link Field Descriptions

| Field | Description |
|---------|-------------|
| Link ID | Link type. |

| Field | Description |
|------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 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. 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/RP0/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 80: 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). |
| 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. |

| Field | Description |
|---------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 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. 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 XR EXEC mode.

show ospf message-queue

This command has no arguments or keywords.

Command Default No default behavior or values

Command Modes XR EXEC

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | 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 | Operations |
|---------|---------|------------|
| | ospf | read |

Examples

The following is sample output from the **show ospf message-queue** command:

```
RP/0/RP0/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 81: 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 enqueueing or dequeuing errors seen across all the linked-lists in the OSPF process. |

Related Commands

| Command | Description |
|--------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------|
| queue dispatch incoming, on page 823 | Limits the number of continuous incoming events processed. |
| queue dispatch rate-limited-lsa, on page 825 | Sets the maximum number of rate-limited link-state advertisements (LSAs) processed per run. |
| queue dispatch spf-lsa-limit, on page 827 | 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 828 | 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 XR EXEC mode.

```
show ospf [process-name] [area-id] neighbor [{type interface-path-id] [neighbor-id] [detail] | area-sorted}]
```

| Syntax Description | |
|--------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>process-name</i> | (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. |
| <i>area-id</i> | (Optional) Area ID. If you do not specify an area, all areas are displayed. |
| <i>type</i> | Interface type. For more information, use the question mark (?) online help function. |
| <i>interface-path-id</i> | 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. |
| <i>neighbor-id</i> | (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. |

| | |
|------------------------|---------------|
| Command Default | All neighbors |
|------------------------|---------------|

| | |
|----------------------|---------|
| Command Modes | XR EXEC |
|----------------------|---------|

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | 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 | 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/RP0/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 82: 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/RP0/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 retrns during last exchange 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 exchange 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 83: 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. |
| 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/RP0/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 exchange 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 84: 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. |
| 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/RP0/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 85: 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/RP0/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 exchange 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 86: 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. |
| 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. |

| Field | Description |
|-----------------------|--------------------------------------------------------------------------------------------------------------------------------|
| 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. |

Related Commands

| Command | Description |
|------------------------------------------|-------------------------------------|
| router ospf, on page 842 | 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 XR EXEC mode.

```
show ospf [process-name] [area-id] request-list [type interface-path-id] [neighbor-id]
```

Syntax Description

| | |
|--------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>process-name</i> | (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. |
| <i>area-id</i> | (Optional) Area ID. If you do not specify an area, all areas are displayed. |
| <i>type</i> | Interface type. For more information, use the question mark (?) online help function. |
| <i>interface-path-id</i> | 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. |
| <i>neighbor-id</i> | (Optional) IP address of the OSPF neighbor. |

Command Default

All neighbors

Command Modes

XR EXEC

Command History

| Release | Modification |
|---------------|------------------------------|
| 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.

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/RP0/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 87: 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 , on page 842 | Configures an OSPF routing process. |
| show ospf retransmission-list , on page 872 | 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 XR EXEC mode.

```
show ospf [process-name] [area-id] retransmission-list [type interface-path-id] [neighbor-id]
```

Syntax Description

| | |
|--------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>process-name</i> | (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. |
| <i>area-id</i> | (Optional) Area ID. If you do not specify an area, all areas are displayed. |
| <i>type</i> | Interface type. For more information, use the question mark (?) online help function. |
| <i>interface-path-id</i> | 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. |
| <i>neighbor-id</i> | (Optional) IP address of the OSPF neighbor. |

Command Default

All neighbors

Command Modes

XR EXEC

Command History

| Release | Modification |
|---------------|------------------------------|
| 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.

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/RP0/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 88: 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, on page 842 | Configures an OSPF routing process. |
| show ospf request-list, on page 870 | 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 XR EXEC mode.

```
show ospf [process-name] routes [{connected | external | local}] [prefix mask] [prefix/length]
```

Syntax Description

| | |
|---------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>process-name</i> | (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. |
| 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). |
| <i>prefix</i> | (Optional) IP prefix, which limits output to a specific route. If the <i>prefix</i> argument is specified, either the <i>length</i> or <i>mask</i> argument is required. |
| <i>mask</i> | (Optional) IP address mask. |
| <i>/ length</i> | (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 <i>length</i> is used, the slash is required. |

Command Default

All route types

Command Modes

XR EXEC

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 5.0.0 | This command was introduced. |

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/RP0/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, 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
    Attributes: Metric: 104, SRLG Disjoint
```

Examples

The following is sample output from the **show ospf routes** command:

```
RP/0/RP0/CPU0:router# show ospf routes

Topology Table for ospf 1 with ID 10.3.4.2

Codes:O - Intra area, 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
```

The following is sample output from the **show ospf routes** command with a process name of 100:

```
RP/0/RP0/CPU0:router# show ospf 100 routes

Topology Table for ospf 100 with ID 172.23.54.14

Codes:O - Intra area, 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 89: 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/RP0/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 90: 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, on page 842 | Configures an OSPF routing process. |
| show route, on page 1101 | Displays current routes information in the Routing Information Base (RIB). |
| show rib opaques, on page 1090 | 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 XR EXEC mode.

show ospf [*process-name*] **sham-links**

| | | |
|---------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------|
| Syntax Description | <i>process-name</i> (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. | |
| Command Default | No default behavior or values | |
| Command Modes | XR EXEC | |
| Command History | Release | Modification |
| | Release 5.0.0 | This command was introduced. |
| Usage Guidelines | <p>To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.</p> <p>Use the show ospf sham-links command to display OSPF sham-link information.</p> | |
| Task ID | Task ID | Operations |
| | ospf | read |

show ospf statistics interface

To display the per interface statistics for OSPFv2, use the **show ospf statistics interface** command in XR EXEC mode.

```
show ospf [process name [area id]] [area id] statistics interface [{interface name | summary-only}]
```

| Syntax Description | |
|---------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>process-name</i> | (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. |
| <i>area id</i> | (Optional) Area number used to define the particular area. |
| summary-only | (Optional) Displays only the summary statistics for the given instance or area (if specified). |

Command Default No default behavior or values.

Command Modes XR EXEC

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | 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 | Operations |
|---------|---------|------------|
| | ospf | read |

Examples

The following is sample output from the **show ospf statistics interface** command:

```
RP/0/RP0/CPU0:router# show ospf 0 1.1.1.1 statistics interface
```

```
Interface POS0/3/0/0 Process ID 0 Area 1.1.1.1
Multi-Adjacency Interface
```

```
OSPF packet and LSA statistics
      RX(hello)  RX(router)      TX      LSA RX      LSA TX
Hello          32             -        33          -          -
DB Des         3             3         2           2           4
LS Req         0             0         1           0           0
LS Upd         5             5         3          18          10
LS Ack         1             1         3          10          18
TOTAL          41             9        42          30          32
```

```
OSPF Header Errors
```

show ospf statistics interface

```

Version          0          LLS          0
Type            0          Auth RX     0
Length          0          Auth TX     0
Checksum        0

OSPF LSA Errors

Type            0          Checksum    0
Length          0          Data        0

OSPF Errors

Bad Source      0          Area Mismatch  0
No Virtual Link 0          Self Originated 0
No Sham Link    0          Duplicate ID    0
Nbr ignored     0          Graceful Shutdown 0
Unknown nbr     0          Passive intf    0
No DR/BDR       0          Disabled intf   0
Enqueue         0          Unspecified RX  0
Socket          0          Unspecified TX  0

```

This table describes the significant fields shown in the display.

Table 91: show ospf statistics interface Field Descriptions

| Field | Description |
|--------------------------------|----------------------------------------------------------------------------------------------|
| OSPF packet and LSA statistics | Packets and LSAs received and transmitted on a given interface. |
| OSPF Header Errors | OSPF packets discarded due to the error in the OSPF header. |
| OSPF LSA Errors | OSPF LSAs discarded due to the error in the OSPF LSA header. |
| OSPF Errors | Packets discarded or errors encountered during handling OSPF packets on the given interface. |

Related Commands

| Command | Description |
|--------------------------------------------------------------|----------------------------------------------------------------------|
| clear ospf statistics interface, on page 739 | Clears the Open Shortest Path First (OSPF) statistics per interface. |

show ospf summary-prefix

To display Open Shortest Path First (OSPF) aggregated summary address information, use the **show ospf summary-prefix** command in XR EXEC mode.

```
show ospf [process-name] summary-prefix
```

| Syntax Description | <i>process-name</i> (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. | | | | |
|---------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------|--------------|---------------|------------------------------|
| Command Default | All summary prefixes | | | | |
| Command Modes | XR EXEC | | | | |
| Command History | <table border="1"> <thead> <tr> <th>Release</th> <th>Modification</th> </tr> </thead> <tbody> <tr> <td>Release 5.0.0</td> <td>This command was introduced.</td> </tr> </tbody> </table> | Release | Modification | Release 5.0.0 | This command was introduced. |
| Release | Modification | | | | |
| Release 5.0.0 | This command was introduced. | | | | |
| Usage Guidelines | <p>To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.</p> <p>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.</p> | | | | |
| Task ID | <table border="1"> <thead> <tr> <th>Task ID</th> <th>Operations</th> </tr> </thead> <tbody> <tr> <td>ospf</td> <td>read</td> </tr> </tbody> </table> | Task ID | Operations | ospf | read |
| Task ID | Operations | | | | |
| ospf | read | | | | |

Examples

The following is sample output from the **show ospf summary-prefix** command:

```
RP/0/RP0/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 92: 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. |

| Field | Description |
|--------|--------------------------------------------------------------------------------------------------|
| 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, on page 842 | Configures an OSPF routing process. |
| summary-prefix (OSPF), on page 893 | 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 XR EXEC mode.

```
show ospf [process-name] virtual-links
```

| | | |
|---------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------|
| Syntax Description | <i>process-name</i> (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. | |
| Command Default | All virtual links | |
| Command Modes | XR EXEC | |
| Command History | Release | Modification |
| | Release 5.0.0 | This command was introduced. |
| Usage Guidelines | <p>To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.</p> <p>Use the show ospf virtual-links command to display useful information for debugging OSPF routing operations.</p> | |
| Task ID | Task ID | Operations |
| | ospf | read |

Examples

The following is sample output from the **show ospf virtual-links** command:

```
RP/0/RP0/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 93: 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 using using 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 , on page 842 | 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 XR EXEC mode.

```
show protocols [{afi-all | ipv4 | ipv6}] [{allprotocol}]
```

| 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. |
| <i>protocol</i> | (Optional) Specifies a routing protocol. For the IPv4 address family, the options are: <ul style="list-style-type: none"> • bgp • eigrp • isis • ospf • rip For the IPv6 address family, the options are: <ul style="list-style-type: none"> • bgp • eigrp • isis • ospfv3 |

Command Default No default behavior or value

Command Modes XR EXEC

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | 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 | Operations |
|---------|---------|------------|
| | ospf | read |
| | rib | read |

Examples

The following is an OSPF configuration and the resulting `show protocols ospf` display:

```
RP/0/RP0/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/RP0/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 94: 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. |

| Field | Description |
|----------------|------------------------------------------------------------------------------------------------------------------------------------------------|
| 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 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

| | |
|----------------------|------------------------------------------------------------|
| <i>window-size</i> | Specifies the trap rate limit sliding window size. |
| <i>max-num-traps</i> | Specifies the maximum number of traps sent in window time. |

Command Default

None

Command Modes

Router configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| 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 |
|---------|------------|
| 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/RP0/CPU0:router (config) #router ospf 100
RP/0/RP0/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 | <p>route-policy <i>policy-name</i> Specifies the route policy to apply to OSPFv2 prefix prioritization.</p> <p>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.</p> | | | | |
|---------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------|--------------|---------------|------------------------------|
| Command Default | SPF prefix prioritization is disabled. | | | | |
| Command Modes | OSPF router configuration | | | | |
| Command History | <table border="1"> <thead> <tr> <th>Release</th> <th>Modification</th> </tr> </thead> <tbody> <tr> <td>Release 5.0.0</td> <td>This command was introduced.</td> </tr> </tbody> </table> | Release | Modification | Release 5.0.0 | This command was introduced. |
| Release | Modification | | | | |
| 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.

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/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# prefix-set ospf-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 ospf-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)# router ospf 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 (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 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.

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/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router ospf 201
RP/0/RP0/CPU0:router(config-ospf)# area 10.0.0.0
RP/0/RP0/CPU0:router(config-ospf-ar)# stub
RP/0/RP0/CPU0:router(config-ospf-ar)# default-cost 20
RP/0/RP0/CPU0:router(config-ospf-ar)# interface GigabitEthernet 0/3/0/3
```

Related Commands

| Command | Description |
|----------------------------------------------------|-----------------------------------------------------------------------|
| authentication (OSPF), on page 725 | Enables authentication for an OSPF area. |
| default-cost (OSPF), on page 747 | 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 | <i>address</i> | Summary address designated for a range of addresses. |
| | <i>mask</i> | 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 <i>tag</i> | (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 | |
| Command History | Release | Modification |
| | 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. | |
| | 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/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router ospf 201
RP/0/RP0/CPU0:router(config-ospf)# summary-prefix 10.1.0.0 255.255.0.0
```

Related Commands

| Command | Description |
|-------------------------------------------|---------------------------------------------------------|
| range (OSPF), on page 830 | 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 | <i>seconds</i> Interval (in seconds) at which LSAs are grouped and refreshed, checksummed, or aged. Range is 10 seconds to 1800 seconds. | | | | |
|---------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------|--------------|---------------|------------------------------|
| Command Default | <i>seconds</i> : 240 seconds | | | | |
| Command Modes | Router configuration | | | | |
| Command History | <table border="1"> <thead> <tr> <th>Release</th> <th>Modification</th> </tr> </thead> <tbody> <tr> <td>Release 5.0.0</td> <td>This command was introduced.</td> </tr> </tbody> </table> | Release | Modification | Release 5.0.0 | This command was introduced. |
| Release | Modification | | | | |
| 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.

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/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router ospf 1
RP/0/RP0/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 | <i>milliseconds</i> Minimum interval (in milliseconds) between accepting same LSA. Range is 0 to 600000 milliseconds. |
|---------------------------|--------------------------------------------------------------------------------------------------------------------------|

| | |
|------------------------|----------------------------------------|
| Command Default | <i>milliseconds</i> : 100 milliseconds |
|------------------------|----------------------------------------|

| | |
|----------------------|----------------------|
| Command Modes | Router configuration |
|----------------------|----------------------|

| Command History | Release | Modification |
|------------------------|----------------|------------------------------|
| | 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 | 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/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router ospf 1
RP/0/RP0/CPU0:router(config-ospf)# timers lsa min-arrival 2
```

timers lsa refresh

To configure the time interval at which Open Shortest Path First (OSPF) self-originated link-state advertisements (LSAs) are refreshed, use the **timers lsa refresh** command in an appropriate configuration mode. To restore the default value, use the **no** form of this command.

timers lsa refresh *seconds*
no timers lsa refresh

| Syntax Description | <i>seconds</i> How often self-originated LSAs should be refreshed, in seconds. Range is 1800 to 2700 seconds. | | | | |
|---------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------|--------------|---------------|------------------------------|
| Command Default | <i>seconds</i> : 1800 seconds. | | | | |
| Command Modes | Router configuration | | | | |
| Command History | <table border="1"> <thead> <tr> <th>Release</th> <th>Modification</th> </tr> </thead> <tbody> <tr> <td>Release 5.0.0</td> <td>This command was introduced.</td> </tr> </tbody> </table> | Release | Modification | Release 5.0.0 | This command was introduced. |
| Release | Modification | | | | |
| 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.

timers lsa refresh command allows self-originated LSAs to be refreshed at non-standard times, anywhere from 1800 to 2700 seconds. Higher refresh interval value may gradually lead to lower CPU utilization by OSPF process.

| Task ID | Task ID | Operations |
|---------|---------|----------------|
| | ospf | read, write |

Examples

The following example shows how to configure an LSA refresh interval of 1800 seconds:

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router ospf 100
RP/0/RP0/CPU0:router(config-ospf)# timers lsa refresh 1800
```

Related Commands

| Command | Description |
|------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------|
| timers lsa group-pacing, on page 895 | Change the interval at which Open Shortest Path First (OSPF) link-state advertisements (LSAs) are collected into a group and refreshed, checksummed, or aged. |

| Command | Description |
|-----------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------|
| timers lsa min-arrival, on page 896 | Limits the frequency that new instances of any particular Open Shortest Path First (OSPF) link-state advertisements (LSAs) can be accepted during flooding. |

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 | |
|-----------------------|----------------------------------------------------------------------------------------------------|
| <i>start-interval</i> | Delay to generate first occurrence of LSA in milliseconds. Range is 0 to 600000 milliseconds. |
| <i>hold-interval</i> | Minimum delay between originating the same LSA in milliseconds. Range is 1 to 600000 milliseconds. |
| <i>max-interval</i> | Maximum delay between originating the same LSA in milliseconds. Range is 1 to 600000 milliseconds. |

| Command Default | |
|-----------------------|---------------------|
| <i>start-interval</i> | : 50 milliseconds |
| <i>hold-interval</i> | : 200 milliseconds |
| <i>max-interval</i> | : 5000 milliseconds |

| Command Modes | |
|---------------|----------------------|
| | Router configuration |

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | 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.

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/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router ospf 1
RP/0/RP0/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/RP0/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 msec
  Minimum hold time between two consecutive SPF's 10000 msec
  Maximum wait time between two consecutive SPF's 10000 msec
  Initial LSA throttle delay 500 msec
  Minimum hold time for LSA throttle 1000 msec
  Maximum wait time for LSA throttle 90000 msec
  Minimum LSA interval 1000 msec. Minimum LSA arrival 1 sec
  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 846 | 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 | | |
|--------------------|---------------------|------------------------------------------------------------------------------------------------------------------|
| | <i>spf-start</i> | Initial SPF schedule delay (in milliseconds). Range is 1 to 600000 milliseconds. |
| | <i>spf-hold</i> | Minimum hold time (in milliseconds) between two consecutive SPF calculations. Range is 1 to 600000 milliseconds. |
| | <i>spf-max-wait</i> | 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 |

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | 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.

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/RP0/CPU0:router# configure  
RP/0/RP0/CPU0:router(config)# router ospf 1  
RP/0/RP0/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 | <i>seconds</i> Time (in seconds) required to send a link-state update. Range is 1 to 65535 seconds. |
|---------------------------|-----------------------------------------------------------------------------------------------------|

| | |
|------------------------|---------------------------|
| Command Default | <i>seconds</i> : 1 second |
|------------------------|---------------------------|

| | |
|----------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Command Modes | Router configuration Area configuration Interface configuration Virtual-link configuration Multi-area configuration Sham-link configuration |
|----------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------|

| Command History | Release | Modification |
|------------------------|----------------|------------------------------|
| | 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. |
|-------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

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/RP0/CPU0:router(config)# router ospf 1
```

```
RP/0/RP0/CPU0:router(config-ospf)# area 0
RP/0/RP0/CPU0:router(config-ospf-ar)# interface GigabitEthernet 0/3/0/0
RP/0/RP0/CPU0:router(config-ospf-ar-if)# transmit-delay 3
```

Related Commands

| Command | Description |
|----------------------------------------|------------------------------------------------------------|
| show ospf, on page 846 | Displays general information about OSPF routing processes. |

ucmp (OSPFv2)

To enable unequal cost multipath (UCMP) calculation for Open Shortest Path First version 2 (OSPFv2), use the **ucmp** command in an appropriate OSPF configuration mode. To disable UCMP for OSPFv2, 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. |
| <i>prefix-list-name</i> | Name of the prefix-list to be specified to filter UCMP paths. |
| variance | (Optional) Specifies variance parameter to filter UCMP paths based on cost. |
| <i>variance-value</i> | Variance value. The variance value is expressed in terms of percentage of the Primary path metric. Range is from 101 to 10000. Default variance value is 200. |

Command Default UCMP is disabled.

Command Modes Router configuration
 VRF configuration

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | Release 5.2.1 | 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.

Enabling the **ucmp** command makes the router calculate UCMP nexthops for all 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 the variance value is 150 and the primary path metric is 100, then select all the UCMP nexthops with metrics from 101 to 150.

| Task ID | Task ID | Operation |
|---------|---------|----------------|
| | ospf | read, write |

This example shows how to enable calculation of UCMP nexthops for all the prefixes in the routing table:

```
RP/0/RP0/CPU0:router#configure
RP/0/RP0/CPU0:router(config)#router ospf 1
RP/0/RP0/CPU0:router(config-ospf)#ucmp
```

This example shows how to enable calculation of UCMP nexthops for a set of prefixes in the prefix-list *list1*:

```
RP/0/RP0/CPU0:router#configure
RP/0/RP0/CPU0:router(config)#router ospf 1
RP/0/RP0/CPU0:router(config-ospf)#ucmp prefix-list list1
```

This example shows how to enable calculation of UCMP mexthops with variance value *120*:

```
RP/0/RP0/CPU0:router#configure
RP/0/RP0/CPU0:router(config)#router ospf 1
RP/0/RP0/CPU0:router(config-ospf)#ucmp variance 120
```

Related Commands

| Command | Description |
|--------------------------------------------------------------|-------------------------------------------------------------------------------|
| ucmp delay-interval (OSPFv2), on page 908 | Specifies delay between primary SPF completion and start of UCMP computation. |
| ucmp exclude interface (OSPFv2), on page 909 | Excludes an interface from unequal cost multipath (UCMP) computation. |
| bandwidth | Configures the bandwidth of an interface. |

ucmp delay-interval (OSPFv2)

To specify delay between primary SPF completion and start of UCMP computation, use the **ucmp delay-interval** command in an appropriate OSPF configuration mode. To disable this functionality, use the **no** form of this command.

ucmp delay-interval *delay-interval*
no ucmp delay-interval

| | |
|---------------------------|-------------------------------------------------------------------------------------------------------------------------------------|
| Syntax Description | <i>delay-interval</i> Delay interval value in milliseconds. Range is from 100 to 65535 . The default value for the interval is 100. |
|---------------------------|-------------------------------------------------------------------------------------------------------------------------------------|

Command Default UCMP delay interval is set to 100 milliseconds.

Command Modes Router configuration
 VRF configuration

| Command History | Release | Modification |
|------------------------|----------------|------------------------------|
| | Release 5.2.1 | 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 delay between primary SPF completion and the start of UCMP computation, to 800 milliseconds:

```
RP/0/RP0/CPU0:router#configure
RP/0/RP0/CPU0:router (config)#router ospf 1
RP/0/RP0/CPU0:router (config-ospf)#ucmp delay-interval 800
```

| Related Commands | Command | Description |
|-------------------------|--------------------------------------------------------------|-----------------------------------------------------------------------|
| | ucmp (OSPFv2), on page 906 | Enables unequal cost multipath (UCMP) calculation for OSPFv2. |
| | ucmp exclude interface (OSPFv2), on page 909 | Excludes an interface from unequal cost multipath (UCMP) computation. |
| | bandwidth | Configures the bandwidth of an interface. |

ucmp exclude interface (OSPFv2)

To exclude an interface from unequal cost multipath (UCMP) computation, use the **ucmp exclude interface** command in an appropriate OSPF 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 | | |
|--------------------|--------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| | <i>type</i> | Interface type. For more information, use the question mark (?) online help function. |
| | <i>interface-path-id</i> | Physical interface or virtual interface. |
| | | <p>Note Use the show interfaces command to see a list of all interfaces currently configured on the router.</p> <p>For more information about the syntax for the router, use the question mark (?) online help function.</p> |

Command Default None

Command Modes Router configuration
VRF configuration

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | Release 5.2.1 | 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 **ucmp exclude interface** command to exclude an interface from being selected as a UCMP backup path.

| Task ID | Task ID | Operation |
|---------|---------|----------------|
| | ospf | read, write |

This example shows how to exclude interface POS 0/3/0/1 from UCMP computation:

```
RP/0/RP0/CPU0:router#configure
RP/0/RP0/CPU0:router(config)#router ospf 1
RP/0/RP0/CPU0:router(config-ospf)#ucmp exclude interface GigabitEthernet 0/3/0/1
```

Related Commands

| Command | Description |
|-----------------------------------------------------------|-------------------------------------------------------------------------------|
| ucmp (OSPFv2), on page 906 | Enables unequal cost multipath (UCMP) calculation for OSPFv2. |
| ucmp delay-interval (OSPFv2), on page 908 | Specifies delay between primary SPF completion and start of UCMP computation. |
| bandwidth | Configures the bandwidth of an interface. |

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 | <i>router-id</i> 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 | <table border="1"> <thead> <tr> <th>Release</th> <th>Modification</th> </tr> </thead> <tbody> <tr> <td>Release 5.0.0</td> <td>This command was introduced.</td> </tr> </tbody> </table> | Release | Modification | Release 5.0.0 | This command was introduced. |
| Release | Modification | | | | |
| 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.

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/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router (config)# router ospf 201
RP/0/RP0/CPU0:router (config-ospf)# area 10.0.0.0
RP/0/RP0/CPU0:router (config-ospf-ar)# virtual-link 10.3.4.5
RP/0/RP0/CPU0:router (config-ospf-ar-vl)#
```

The following example shows how to establish a virtual link with clear text authentication called mykey:

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router ospf 201
RP/0/RP0/CPU0:router(config-ospf)# area 10.0.0.0
RP/0/RP0/CPU0:router(config-ospf-ar)# virtual-link 10.3.4.5
RP/0/RP0/CPU0:router(config-ospf-ar-vl)# authentication-key 0 mykey
```

Related Commands

| Command | Description |
|------------------------------------------------------|-----------------------------------------------------------------|
| authentication (OSPF), on page 725 | Enables authentication for an OSPF area. |
| show ospf virtual-links, on page 883 | Displays parameters and the current state of OSPF virtual links |



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 Software* module in the *Routing Configuration Guide for Cisco NCS 6000 Series Routers*.

- [address-family \(OSPFv3\)](#), on page 915
- [area \(OSPFv3\)](#), on page 916
- [authentication \(OSPFv3\)](#), on page 917
- [auto-cost \(OSPFv3\)](#), on page 919
- [clear ospfv3 process](#), on page 921
- [clear ospfv3 redistribution](#), on page 923
- [clear ospfv3 routes](#), on page 924
- [clear ospfv3 statistics](#), on page 925
- [cost \(OSPFv3\)](#), on page 927
- [database-filter all out \(OSPFv3\)](#), on page 929
- [dead-interval \(OSPFv3\)](#), on page 931
- [default-cost \(OSPFv3\)](#), on page 933
- [default-information originate \(OSPFv3\)](#), on page 935
- [default-metric \(OSPFv3\)](#), on page 937
- [demand-circuit \(OSPFv3\)](#), on page 939
- [distance ospfv3](#), on page 941
- [distribute-list prefix-list in](#), on page 943
- [distribute-list prefix-list out](#), on page 945
- [encryption](#), on page 947
- [flood-reduction \(OSPFv3\)](#), on page 949
- [graceful-restart \(OSPFv3\)](#), on page 951
- [hello-interval \(OSPFv3\)](#), on page 952
- [instance \(OSPFv3\)](#), on page 954
- [interface \(OSPFv3\)](#), on page 956
- [log adjacency changes \(OSPFv3\)](#), on page 958
- [maximum interfaces \(OSPFv3\)](#), on page 959
- [maximum paths \(OSPFv3\)](#), on page 960
- [maximum redistributed-prefixes \(OSPFv3\)](#), on page 961
- [mtu-ignore \(OSPFv3\)](#), on page 963

- neighbor (OSPFv3), on page 965
- network (OSPFv3), on page 967
- nssa (OSPFv3), on page 969
- nsr (OSPFv3), on page 971
- ospfv3 name-lookup, on page 973
- packet-size (OSPFv3), on page 974
- passive (OSPFv3), on page 975
- priority (OSPFv3), on page 977
- protocol shutdown (OSPFv3), on page 979
- range (OSPFv3), on page 980
- redistribute (OSPFv3), on page 982
- retransmit-interval (OSPFv3), on page 986
- router-id (OSPFv3), on page 988
- router ospfv3, on page 990
- show ospfv3, on page 991
- show ospfv3 border-routers, on page 995
- show ospfv3 database, on page 997
- show ospfv3 flood-list, on page 1009
- show ospfv3 interface, on page 1011
- show ospfv3 message-queue, on page 1014
- show ospfv3 neighbor, on page 1016
- show ospfv3 request-list, on page 1022
- show ospfv3 retransmission-list, on page 1024
- show ospfv3 routes, on page 1026
- show ospfv3 summary-prefix, on page 1028
- show ospfv3 virtual-links, on page 1030
- show protocols (OSPFv3), on page 1032
- stub (OSPFv3), on page 1034
- stub-router, on page 1036
- summary-prefix (OSPFv3), on page 1038
- timers lsa arrival, on page 1040
- timers pacing flood, on page 1041
- timers pacing lsa-group, on page 1043
- timers pacing retransmission , on page 1045
- timers throttle lsa all (OSPFv3), on page 1047
- timers throttle spf (OSPFv3), on page 1049
- trace (OSPFv3), on page 1051
- transmit-delay (OSPFv3), on page 1053
- virtual-link (OSPFv3), on page 1055

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 ospfv3 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 | <p>ipv6 Specifies IP Version 6 (IPv6) address prefixes.</p> <p>unicast (Optional) Specifies unicast address prefixes.</p> | | | | |
|---------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------|--------------|---------------|------------------------------|
| Command Default | An address family is not specified. | | | | |
| Command Modes | Router ospfv3 configuration | | | | |
| Command History | <table border="1"> <thead> <tr> <th>Release</th> <th>Modification</th> </tr> </thead> <tbody> <tr> <td>Release 5.0.0</td> <td>This command was introduced.</td> </tr> </tbody> </table> | Release | Modification | Release 5.0.0 | This command was introduced. |
| Release | Modification | | | | |
| 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 | <table border="1"> <thead> <tr> <th>Task ID</th> <th>Operations</th> </tr> </thead> <tbody> <tr> <td>ospf</td> <td>read, write</td> </tr> </tbody> </table> | Task ID | Operations | ospf | read, write |
| 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/RP0/CPU0:router(config)# router ospfv3 1
RP/0/RP0/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 | <i>area-id</i> Identifier of an OSPFv3 area. The <i>area-id</i> 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 | <table border="1"> <thead> <tr> <th>Release</th> <th>Modification</th> </tr> </thead> <tbody> <tr> <td>Release 5.0.0</td> <td>This command was introduced.</td> </tr> </tbody> </table> | Release | Modification | Release 5.0.0 | This command was introduced. |
| Release | Modification | | | | |
| 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.

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/RP0/CPU0:router(config)# router ospfv3 1
RP/0/RP0/CPU0:router(config-ospfv3)# area 0
RP/0/RP0/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. | |
| <i>password</i> | 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 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. |

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**

- 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/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router (config)# router ospfv3 201
RP/0/RP0/CPU0:router (config-ospfv3)# router-id 10.1.1.1
RP/0/RP0/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 <i>mbps</i> (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 | <i>mbps</i> : 100 Mbps | |
| Command Modes | Router ospfv3 configuration | |
| Command History | Release | Modification |
| | 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.

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/RP0/CPU0:router(config)# router ospfv3 1  
RP/0/RP0/CPU0:router(config-ospfv3)# auto-cost reference-bandwidth 64
```

Related Commands

| Command | Description |
|--------------------------------------------|-------------------------------------------------------------------------------------|
| cost (OSPFv3), on page 927 | Explicitly specifies the cost of the interface (network) for OSPF path calculation. |

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 XR EXEC mode.

clear ospfv3 [*process-name*] **process**

| | |
|---------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax Description | <i>process-name</i> (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 | XR EXEC |
|----------------------|---------|

| | | |
|------------------------|----------------|------------------------------|
| Command History | Release | Modification |
| | 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.

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 988](#) command.

| Task ID | Task ID | Operations |
|----------------|----------------|-------------------|
| | ospf | read, write |

Examples

The following example shows how to reset all OSPFv3 processes:

```
RP/0/RP0/CPU0:router# clear ospfv3 process
```

The following example shows how to reset the OSPFv3 process 1:

```
RP/0/RP0/CPU0:router# clear ospfv3 1 process
```

Related Commands

| Command | Description |
|-------------------------------------------------|------------------------------------------------|
| router-id (OSPFv3), on page 988 | 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 XR EXEC mode.

clear ospfv3 [*process-name*] **redistribution**

| Syntax Description | <i>process-name</i> (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 | XR EXEC | | | | |
| Command History | <table border="1"> <thead> <tr> <th>Release</th> <th>Modification</th> </tr> </thead> <tbody> <tr> <td>Release 5.0.0</td> <td>This command was introduced.</td> </tr> </tbody> </table> | Release | Modification | Release 5.0.0 | This command was introduced. |
| Release | Modification | | | | |
| 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.

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/RP0/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 XR EXEC mode.

clear ospfv3 [*process-name*] **routes**

| | |
|---------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax Description | <i>process-name</i> (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 | XR EXEC |
|----------------------|---------|

| Command History | Release | Modification |
|------------------------|----------------|------------------------------|
| | 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.

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/RP0/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 XR EXEC mode.

```
clear ospfv3 [process-name] statistics [neighbor [type interface-path-id] [router-id]]
```

| Syntax Description | |
|--------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>process-name</i> | (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. |
| <i>type</i> | Interface type. For more information, use the question mark (?) online help function. |
| <i>interface-path-id</i> | 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. |
| <i>router-id</i> | (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 XR EXEC

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | 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.

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/RP0/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 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.

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:

$$10^8 / \text{bandwidth}$$

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/RP0/CPU0:router(config)# router ospfv3 201  
RP/0/RP0/CPU0:router(config-ospfv3)# area 0  
RP/0/RP0/CPU0:router(config-ospfv3-ar)# interface GigabitEthernet 0/1/0/1  
RP/0/RP0/CPU0:router(config-ospfv3-ar-if)# cost 65
```

Related Commands

| Command | Description |
|-------------------------------------------------|-------------------------------------------------------------------------------|
| auto-cost (OSPFv3), on page 919 | 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 | <p>If this command is not specified in interface configuration mode, then the interface adopts the database filter parameter specified for the area.</p> <p>If this command is not specified in area configuration mode, then the interface adopts the database filter parameter specified for the process.</p> <p>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.</p> |

| | |
|----------------------|---------------------------------------------------------------------------------------------|
| Command Modes | <p>Interface configuration</p> <p>Area configuration</p> <p>Router OSPFv3 configuration</p> |
|----------------------|---------------------------------------------------------------------------------------------|

| | | |
|------------------------|----------------|------------------------------|
| Command History | Release | Modification |
| | Release 5.0.0 | This command was introduced. |

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

| | | |
|----------------|----------------|-------------------|
| 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/RP0/CPU0:router(config)# router ospfv3 1
RP/0/RP0/CPU0:router(config-ospfv3)# area 0
RP/0/RP0/CPU0:router(config-ospfv3-ar)# interface GigabitEthernet 0/2/0/3
RP/0/RP0/CPU0:router(config-ospfv3-ar-if)# database-filter all out
```

Related Commands

| Command | Description |
|------------------------------------------------|---------------------------------------------------------------------|
| neighbor (OSPFv3), on page 965 | 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 | <i>seconds</i> 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 | <p>If this command is not specified in interface configuration mode, then the interface adopts the dead interval parameter specified for the area.</p> <p>If this command is not specified in area configuration mode, then the interface adopts the dead interval parameter specified for the process.</p> <p>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.</p> | | | | |
| Command Modes | <p>Interface configuration</p> <p>Area configuration</p> <p>Router OSPFv3 configuration</p> <p>Virtual-link configuration</p> | | | | |
| Command History | <table border="1"> <thead> <tr> <th>Release</th> <th>Modification</th> </tr> </thead> <tbody> <tr> <td>Release 5.0.0</td> <td>This command was introduced.</td> </tr> </tbody> </table> | Release | Modification | Release 5.0.0 | This command was introduced. |
| Release | Modification | | | | |
| Release 5.0.0 | This command was introduced. | | | | |
| Usage Guidelines | <p>To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.</p> <p>Two Open Shortest Path First Version 3 (OSPFv3) routers do not become adjacent if their dead interval values differ.</p> <p>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.</p> | | | | |
| Task ID | <table border="1"> <thead> <tr> <th>Task ID</th> <th>Operations</th> </tr> </thead> <tbody> <tr> <td>ospf</td> <td>read, write</td> </tr> </tbody> </table> | Task ID | Operations | ospf | read, write |
| 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/RP0/CPU0:router(config)# router ospfv3 1
RP/0/RP0/CPU0:router(config-ospfv3)# area 0
RP/0/RP0/CPU0:router(config-ospfv3-ar)# interface GigabitEthernet 0/2/0/3
RP/0/RP0/CPU0:router(config-ospfv3-ar-if)# dead-interval 40
```

Related Commands

| Command | Description |
|------------------------------------------------------|-----------------------------------------------------------------------------------------------------|
| hello-interval (OSPFv3), on page 952 | 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 | <i>cost</i> 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 | <i>cost</i> : 1 | |
| Command Modes | Area configuration | |
| Command History | Release | Modification |
| | Release 5.0.0 | This command was introduced. |
| Usage Guidelines | <p>To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.</p> <p>Use the default-cost command only on an Area Border Router (ABR) attached to a stub or an NSSA area.</p> <p>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.</p> | |
| Task ID | Task ID | Operations |
| | ospf | read, write |
| Examples | <p>The following example shows how to assign a cost of 20 to the default route sent into area 10.15.0.0:</p> <pre>RP/0/RP0/CPU0:router(config)# router ospfv3 201 RP/0/RP0/CPU0:router(config-ospfv3)# area 10.15.0.0 RP/0/RP0/CPU0:router(config-ospfv3-ar)# stub RP/0/RP0/CPU0:router(config-ospfv3-ar)# default-cost 20 RP/0/RP0/CPU0:router(config-ospfv3-ar)# interface GigabitEthernet 0/3/0/1</pre> | |

Related Commands

| Command | Description |
|---------------------------------------------|---------------------------------|
| stub (OSPFv3), on page 1034 | 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 <i>policy-name</i> | (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 <i>metric-value</i> | (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 <i>type-value</i> | (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 <i>tag-value</i> | (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 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.

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 NCS 6000 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 NCS 6000 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/RP0/CPU0:router(config)#router ospfv3 109
RP/0/RP0/CPU0:router(config-ospfv3)#default-information originate metric 100 metric-type 1
```

Related Commands

| Command | Description |
|----------------------------------------------------|---------------------------------------------------------------------------|
| redistribute (OSPFv3), on page 982 | 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 | <i>value</i> 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 | <table border="1"> <thead> <tr> <th>Release</th> <th>Modification</th> </tr> </thead> <tbody> <tr> <td>Release 5.0.0</td> <td>This command was introduced.</td> </tr> </tbody> </table> | Release | Modification | Release 5.0.0 | This command was introduced. |
| Release | Modification | | | | |
| 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.

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/RP0/CPU0:router(config)# router ospfv3 1
RP/0/RP0/CPU0:router(config-ospfv3)# default-metric 10
RP/0/RP0/CPU0:router(config-ospfv3)# redistribute isis IS-IS_osp
```

Related Commands

| Command | Description |
|----------------------------------------------------|---------------------------------------------------------------------------|
| redistribute (OSPFv3), on page 982 | 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 | <p>If this command is not specified in interface configuration mode, then the interface adopts the demand circuit parameter specified for the area.</p> <p>If this command is not specified in area configuration mode, then the interface adopts the demand circuit parameter specified for the process.</p> <p>If this command is not specified at any level, then the interface is not a demand circuit.</p> | | | | |
| Command Modes | <p>Interface configuration</p> <p>Area configuration</p> <p>Router OSPFv3 configuration</p> <p>Virtual-link configuration</p> | | | | |
| Command History | <table border="1"> <thead> <tr> <th>Release</th> <th>Modification</th> </tr> </thead> <tbody> <tr> <td>Release 5.0.0</td> <td>This command was introduced.</td> </tr> </tbody> </table> | Release | Modification | Release 5.0.0 | This command was introduced. |
| Release | Modification | | | | |
| Release 5.0.0 | This command was introduced. | | | | |
| Usage Guidelines | <p>To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.</p> <p>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.</p> | | | | |
| Task ID | <table border="1"> <thead> <tr> <th>Task ID</th> <th>Operations</th> </tr> </thead> <tbody> <tr> <td>ospf</td> <td>read, write</td> </tr> </tbody> </table> | Task ID | Operations | ospf | read, write |
| 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/RP0/CPU0:router(config)# router ospfv3 1  
RP/0/RP0/CPU0:router(config-ospfv3)# area 0  
RP/0/RP0/CPU0:router(config-ospfv3)# interface GigabitEthernet 0/3/0/1  
RP/0/RP0/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. |
| | <i>distance</i> | The route administrative distance. |
| Command Default | <i>distance</i> : 110 | |
| Command Modes | Router ospfv3 configuration | |
| Command History | Release | Modification |
| | 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.

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/RP0/CPU0:router(config)# router ospfv3 1
RP/0/RP0/CPU0:router(config-ospfv3)# redistribute ospfv3 2
```

distance ospfv3

```
RP/0/RP0/CPU0:router(config-ospfv3)# distance ospfv3 external 200
RP/0/RP0/CPU0:router(config-ospfv3)# exit
RP/0/RP0/CPU0:router(config)# router ospfv3 2
RP/0/RP0/CPU0:router(config-ospfv3)# redistribute ospfv3 1
RP/0/RP0/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 | <i>prefix-list-name</i> 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 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. |
|-------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

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/RP0/CPU0:router(config)# ipv6 prefix-list preflist1
RP/0/RP0/CPU0:router(config-ipv6-pfx)# deny 2001:e624::/32 le 128
RP/0/RP0/CPU0:router(config-ipv6-pfx)# permit ::/0 le 128
!
RP/0/RP0/CPU0:router(config)# ipv6 prefix-list preflist2
RP/0/RP0/CPU0:router(config-ipv6-pfx)# deny 2002::/16
RP/0/RP0/CPU0:router(config-ipv6-pfx)# permit ::/0 le 128
!
RP/0/RP0/CPU0:router(config)# router ospfv3 1
RP/0/RP0/CPU0:router(config-ospfv3)# distribute-list prefix-list preflist1 in
RP/0/RP0/CPU0:router(config-ospfv3)# area 1
RP/0/RP0/CPU0:router(config-ospfv3-ar)# interface GigabitEthernet 0/2/0/0
RP/0/RP0/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) in 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 982](#) command are redistributed into OSPFv3.

Command Modes

Router OSPFv3 configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| 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.

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/RP0/CPU0:router(config)# ipv6 prefix-list p1
RP/0/RP0/CPU0:router(config-ipv6-px)# deny 2001:e624::/32 le 128
RP/0/RP0/CPU0:router(config-ipv6-px)# permit ::/0 le 128
!
RP/0/RP0/CPU0:router(config)# ipv6 prefix-list p2
RP/0/RP0/CPU0:router(config-ipv6-px)# deny 2064::/16 le 128
RP/0/RP0/CPU0:router(config-ipv6-px)# permit ::/0 le 128
!
RP/0/RP0/CPU0:router(config)# ipv6 prefix-list p3
RP/0/RP0/CPU0:router(config-ipv6-px)# permit 2064::/16 le 128
!
RP/0/RP0/CPU0:router(config)# router ospfv3 1
RP/0/RP0/CPU0:router(config-ospfv3)# redistribute bgp 1
RP/0/RP0/CPU0:router(config-ospfv3)# redistribute bgp 5
RP/0/RP0/CPU0:router(config-ospfv3)# distribute-list prefix-list p1 out
RP/0/RP0/CPU0:router(config-ospfv3)# distribute-list prefix-list p2 out bgp 1
RP/0/RP0/CPU0:router(config-ospfv3)# distribute-list prefix-list p3 out bgp 5
```

Related Commands

| Command | Description |
|----------------------------------------------------|--------------------------------------------------------------------------------------|
| redistribute (OSPFv3), on page 982 | Redistributes routes from one routing domain into another routing domain for OSPFv3. |

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 | Parameter | 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. |
| | <i>spi-value</i> | 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. |
| | <i>encrypt-password</i> | Any contiguous string that can be entered from the keyboard as the encryption password. |
| | <i>auth-password</i> | 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 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.

Use the **encryption** command to encrypt and authenticate OSPFv3 packets.

**Note**

- 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/RP0/CPU0:router(config)#router ospfv3 1
RP/0/RP0/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 | <p>disable (Optional) Turns off this functionality at a specific level.</p> <p>Note The disable keyword is not available in router ospfv3 configuration mode.</p> | | | | |
|---------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------|--------------|---------------|------------------------------|
| Command Default | <p>If this command is not specified in interface configuration mode, then the interface adopts the flood reduction parameter specified by area.</p> <p>If this command is not specified in area configuration mode, then the interface adopts the flood reduction parameter specified for the process.</p> <p>If this command is not specified at any level, then flood reduction is disabled.</p> | | | | |
| Command Modes | <p>Interface configuration</p> <p>Area configuration</p> <p>Router OSPFv3 configuration</p> | | | | |
| Command History | <table border="1"> <thead> <tr> <th>Release</th> <th>Modification</th> </tr> </thead> <tbody> <tr> <td>Release 5.0.0</td> <td>This command was introduced.</td> </tr> </tbody> </table> | Release | Modification | Release 5.0.0 | This command was introduced. |
| Release | Modification | | | | |
| Release 5.0.0 | This command was introduced. | | | | |
| Usage Guidelines | <p>To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.</p> <p>All routers that support Open Shortest Path First Version 3 (OSPFv3) demand circuits are compatible with routers supporting flooding reduction.</p> | | | | |
| Task ID | <table border="1"> <thead> <tr> <th>Task ID</th> <th>Operations</th> </tr> </thead> <tbody> <tr> <td>ospf</td> <td>read, write</td> </tr> </tbody> </table> | Task ID | Operations | ospf | read, write |
| Task ID | Operations | | | | |
| ospf | read, write | | | | |
| Examples | <p>The following example shows how to reduce the flooding of unnecessary LSAs for area 0:</p> <pre>RP/0/RP0/CPU0:router(config)# router ospfv3 1 RP/0/RP0/CPU0:router(config-ospfv3)# area 0</pre> | | | | |

```
RP/0/RP0/CPU0:router(config-ospfv3-ar)# interface GigabitEthernet 0/1/0/3  
RP/0/RP0/CPU0:router(config-ospfv3-ar-if)# flood-reduction
```

Related Commands

| Command | Description |
|-----------------------------------------------------|------------------------------------------------------------------------|
| show ospfv3 interface, on page 1011 | Displays OSPFv3-related interface information. |
| show ospfv3 neighbor, on page 1016 | 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 <i>interval</i> | (Optional) Specifies the minimum interval between graceful restarts. Range is 90 to 3600 seconds. |
| | lifetime <i>lifetime</i> | (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 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 | 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/RP0/CPU0:router(config)# router ospfv3 1
RP/0/RP0/CPU0:router(config-ospfv3)# graceful-restart interval 300
```

| Related Commands | Command | Description |
|------------------|-----------------------------------------------------|------------------------------------------------------------------------|
| | show ospfv3 interface, on page 1011 | Displays OSPFv3-related interface information. |
| | show ospfv3 neighbor, on page 1016 | 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 | <i>seconds</i> Interval (in seconds). The value must be the same for all nodes on a specific network. | |
| Command Default | <p>If this command is not specified in interface configuration mode, then the interface adopts the hello interval parameter specified by area.</p> <p>If this command is not specified in area configuration mode, then the interface adopts the hello interval parameter specified for the process.</p> <p>If this command is not specified at any level, then the hello interval is 10 seconds (broadcast) or 30 seconds (non-broadcast).</p> | |
| Command Modes | <p>Interface configuration</p> <p>Area configuration</p> <p>Router OSPFv3 configuration</p> <p>Virtual-link configuration</p> | |
| Command History | Release | Modification |
| | Release 5.0.0 | This command was introduced. |
| Usage Guidelines | <p>To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.</p> <p>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.</p> | |
| Task ID | Task ID | Operations |
| | ospf | read, write |
| Examples | <p>The following example shows how to set the interval between hello packets to 15 seconds on GigabitEthernet interface 0/3/0/2:</p> | |

```
RP/0/RP0/CPU0:router(config)# router ospfv3 1  
RP/0/RP0/CPU0:router(config-ospfv3)# area 0  
RP/0/RP0/CPU0:router(config-ospfv3-ar)# interface GigabitEthernet 0/3/0/2  
RP/0/RP0/CPU0:router(config-ospfv3-ar-if)# hello-interval 15
```

Related Commands

| Command | Description |
|-----------------------------------------------------|-----------------------------------------------------------------------------------------------|
| dead-interval (OSPFv3), on page 931 | 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 | <i>instance-id</i> 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 | <p>If this command is not specified in interface configuration mode, then the interface adopts the instance parameter specified by the area.</p> <p>If this command is not specified in area configuration mode, then the interface adopts the instance parameter specified for the process.</p> <p>If this command is not specified at any level, then the instance is 0.</p> |
|------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

| | |
|----------------------|---------------------------------------------------------------------------------------------|
| Command Modes | <p>Interface configuration</p> <p>Area configuration</p> <p>Router OSPFv3 configuration</p> |
|----------------------|---------------------------------------------------------------------------------------------|

| | | |
|------------------------|----------------|------------------------------|
| Command History | Release | Modification |
| | Release 5.0.0 | This command was introduced. |

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

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 | <p>The <i>instance-id</i> argument should not be confused with the <i>process-name</i> 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.</p> |
|-------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

| 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/RP0/CPU0:router(config)# router ospfv3 1
RP/0/RP0/CPU0:router(config-ospfv3)# area 0
RP/0/RP0/CPU0:router(config-ospfv3-ar)# interface GigabitEthernet 0/3/0/1
RP/0/RP0/CPU0:router(config-ospfv3-ar-if)# instance 42
```

Related Commands

| Command | Description |
|-------------------------------------------------|--------------------------------------------------------|
| router ospfv3, on page 990 | Configures an OSPFv3 routing process. |
| router-id (OSPFv3), on page 988 | 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 | <i>type</i> | Interface type. For more information, use the question mark (?) online help function. |
| | <i>interface-path-id</i> | 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 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.

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/RP0/CPU0:router(config)# router ospfv3 1
RP/0/RP0/CPU0:router(config-ospfv3)# area 1
RP/0/RP0/CPU0:router(config-ospfv3-ar)# interface GigabitEthernet 0/3/0/1
RP/0/RP0/CPU0:router(config-ospfv3-ar-if)# cost 40
RP/0/RP0/CPU0:router(config-ospfv3-ar-if)# exit
RP/0/RP0/CPU0:router(config-ospfv3-ar)# interface GigabitEthernet 0/3/0/2
RP/0/RP0/CPU0:router(config-ospfv3-ar-if)# cost 65
RP/0/RP0/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 | <p>detail (Optional) Provides all (DOWN, INIT, 2WAY, EXSTART, EXCHANGE, LOADING, FULL) adjacency state changes.</p> <p>disable (Optional) Disables the neighbor state change messages.</p> |
|---------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

Command Default Neighbor state change messages are enabled.

Command Modes Router OSPFv3 configuration

| | | |
|------------------------|----------------|------------------------------|
| Command History | Release | Modification |
| | 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.

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/RP0/CPU0:router(config)# router ospfv3 1
RP/0/RP0/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/RP0/CPU0:router(config)# router ospfv3 1
RP/0/RP0/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 | <i>number-interfaces</i> 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 | <table border="1"> <thead> <tr> <th>Release</th> <th>Modification</th> </tr> </thead> <tbody> <tr> <td>Release 5.0.0</td> <td>This command was introduced.</td> </tr> </tbody> </table> | Release | Modification | Release 5.0.0 | This command was introduced. |
| Release | Modification | | | | |
| 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 | <table border="1"> <thead> <tr> <th>Task ID</th> <th>Operations</th> </tr> </thead> <tbody> <tr> <td>ospf</td> <td>read, write</td> </tr> </tbody> </table> | Task ID | Operations | ospf | read, write |
| Task ID | Operations | | | | |
| ospf | read, write | | | | |
| Examples | <p>This example shows how to allow a maximum of 1500 interfaces in an OSPFv3 process:</p> <pre>RP/0/RP0/CPU0:router(config)# router ospfv3 1 RP/0/RP0/CPU0:router(config-ospfv3)# maximum interfaces 1500</pre> | | | | |

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 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.

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/RP0/CPU0:router(config)# router ospfv3 1
RP/0/RP0/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 | |
|---------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>limit</i> | 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. |
| <i>threshold</i> | (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 <i>limit</i> 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 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 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/RP0/CPU0:router(config)# router ospfv3 1
RP/0/RP0/CPU0:router(config-ospfv3)# redistribute bgp 2406
RP/0/RP0/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 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.

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/RP0/CPU0:router(config)# router ospfv3 1  
RP/0/RP0/CPU0:router(config-ospfv3)# area 0  
RP/0/RP0/CPU0:router(config-ospfv3-ar)# interface GigabitEthernet 0/1/0/3  
RP/0/RP0/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 | | |
|-------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--|
| <i>ipv6-address</i> | 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 <i>number</i> | (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 <i>seconds</i> | (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 <i>number</i> | (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 <i>number</i> : | 0 |
| poll-interval <i>seconds</i> : | 120 seconds (2 minutes) |

| Command Modes | |
|-------------------------|--|
| Interface configuration | |

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | 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.

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/RP0/CPU0:router(config)# router ospfv3 1
RP/0/RP0/CPU0:router(config-ospfv3)# area 0
RP/0/RP0/CPU0:router(config-ospfv3-ar)# interface GigabitEthernet 0/2/0/0
RP/0/RP0/CPU0:router(config-ospfv3-ar)# network non-broadcast
RP/0/RP0/CPU0:router(config-ospfv3-ar-if)# neighbor fe80::3203:a0ff:fe9d:f3fe
```

Related Commands

| Command | Description |
|------------------------------------------------|-----------------------------------------------------------------------------------------|
| priority (OSPFv3), on page 977 | 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 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. | |
| | 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/RP0/CPU0:router(config)# router ospfv3 1
RP/0/RP0/CPU0:router(config-ospfv3)# area 0
RP/0/RP0/CPU0:router(config-ospfv3-ar)# interface TenGigE0/1/0/3
RP/0/RP0/CPU0:router(config-ospfv3-ar-if)# network point-to-point
```

Related Commands

| Command | Description |
|------------------------------------------------|---------------------------------------------------------------------|
| neighbor (OSPFv3), on page 965 | 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 <i>metric-value</i> | (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 <i>type-value</i> | (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 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.

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/RP0/CPU0:router(config)# router ospfv3 1
RP/0/RP0/CPU0:router(config-ospfv3)# router-id 10.18.1.1
RP/0/RP0/CPU0:router(config-ospfv3)# area 1
RP/0/RP0/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 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.

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/RP0/CPU0:router#configure
RP/0/RP0/CPU0:router (config)#router ospfv3 211
RP/0/RP0/CPU0:router (config-ospfv3)#nsr
```

This example shows how to configure NSR for OSPFv3 process 211:

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

```
RP/0/RP0/CPU0:router(config)#router ospfv3 211  
RP/0/RP0/CPU0:router(config-ospfv3)#nsr disable
```

Related Commands

| Command | Description |
|--------------------------------------------|--------------------|
| router ospfv3, on page 990 | |
| show ospfv3, on page 991 | |

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 XR Config 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 XR Config

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | 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.

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/RP0/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 | <i>bytes</i> 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 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. |
|-------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

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/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router (config)# router ospf osp3
RP/0/RP0/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 | <p>disable (Optional) Sends OSPFv3 updates.</p> <p>Note The disable keyword is not available in router ospfv3 configuration mode.</p> | | | | |
|---------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------|--------------|---------------|------------------------------|
| Command Default | <p>If this command is not specified in interface configuration mode, then the interface adopts the passive parameter specified by the area.</p> <p>If this command is not specified in area configuration mode, then the interface adopts the passive parameter specified for the process.</p> <p>If this command is not specified at any level, then the passive parameter is disabled and OSPFv3 updates are sent on the interface.</p> | | | | |
| Command Modes | <p>Interface configuration</p> <p>Area configuration</p> <p>Router OSPFv3 configuration</p> | | | | |
| Command History | <table border="1"> <thead> <tr> <th>Release</th> <th>Modification</th> </tr> </thead> <tbody> <tr> <td>Release 5.0.0</td> <td>This command was introduced.</td> </tr> </tbody> </table> | Release | Modification | Release 5.0.0 | This command was introduced. |
| Release | Modification | | | | |
| Release 5.0.0 | This command was introduced. | | | | |
| Usage Guidelines | <p>To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.</p> <p>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.</p> | | | | |
| Task ID | <table border="1"> <thead> <tr> <th>Task ID</th> <th>Operations</th> </tr> </thead> <tbody> <tr> <td>ospf</td> <td>read, write</td> </tr> </tbody> </table> | Task ID | Operations | ospf | read, write |
| Task ID | Operations | | | | |
| ospf | read, write | | | | |
| Examples | <p>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.</p> | | | | |

```
RP/0/RP0/CPU0:router(config)# router ospfv3 1
RP/0/RP0/CPU0:router(config-ospfv3)# router-id 10.0.0.206
RP/0/RP0/CPU0:router(config-ospfv3)# passive
RP/0/RP0/CPU0:router(config-ospfv3)# area 0
RP/0/RP0/CPU0:router(config-ospfv3-ar)# interface GigabitEthernet 0/3/0/0
RP/0/RP0/CPU0:router(config-ospfv3-ar-if)# passive disable
RP/0/RP0/CPU0:router(config-ospfv3-ar-if)# exit
RP/0/RP0/CPU0:router(config-ospfv3-ar)# interface GigabitEthernet 0/3/0/1
RP/0/RP0/CPU0:router(config-ospfv3-ar-if)# exit
RP/0/RP0/CPU0:router(config-ospfv3-ar)# exit
RP/0/RP0/CPU0:router(config-ospfv3)# area 1
RP/0/RP0/CPU0:router(config-ospfv3-ar)# passive disable
RP/0/RP0/CPU0:router(config-ospfv3-ar)# interface GigabitEthernet 0/2/0/0
RP/0/RP0/CPU0:router(config-ospfv3-ar-if)# exit
RP/0/RP0/CPU0:router(config-ospfv3-ar)# interface GigabitEthernet 0/2/0/1
RP/0/RP0/CPU0:router(config-ospfv3-ar-if)# passive
RP/0/RP0/CPU0:router(config-ospfv3-ar-if)# exit
RP/0/RP0/CPU0:router(config-ospfv3-ar)# interface GigabitEthernet 0/2/0/2
RP/0/RP0/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 | <i>value</i> 8-bit unsigned integer indicating the router priority value. Range is 0 to 255. | | | | |
|---------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------|--------------|---------------|------------------------------|
| Command Default | <p>If this command is not specified in interface configuration mode, then the interface adopts the priority parameter specified by the area.</p> <p>If this command is not specified in area configuration mode, then the interface adopts the priority parameter specified by the process.</p> <p>If this command is not specified at any level, then the default priority is 1.</p> | | | | |
| Command Modes | <p>Interface configuration</p> <p>Area configuration</p> <p>Router OSPFv3 configuration</p> | | | | |
| Command History | <table border="1"> <thead> <tr> <th>Release</th> <th>Modification</th> </tr> </thead> <tbody> <tr> <td>Release 5.0.0</td> <td>This command was introduced.</td> </tr> </tbody> </table> | Release | Modification | Release 5.0.0 | This command was introduced. |
| Release | Modification | | | | |
| Release 5.0.0 | This command was introduced. | | | | |
| Usage Guidelines | <p>To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.</p> <p>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.</p> | | | | |
| Task ID | <table border="1"> <thead> <tr> <th>Task ID</th> <th>Operations</th> </tr> </thead> <tbody> <tr> <td>ospf</td> <td>read, write</td> </tr> </tbody> </table> | Task ID | Operations | ospf | read, write |
| Task ID | Operations | | | | |
| ospf | read, write | | | | |
| Examples | <p>The following example shows how to set the router priority value to 4 on GigabitEthernet interface 0/1/0/1:</p> <pre>RP/0/RP0/CPU0:router(config)# router ospfv3 1</pre> | | | | |

```
RP/0/RP0/CPU0:router(config-ospfv3)# area 0
RP/0/RP0/CPU0:router(config-ospfv3-ar)# interface GigabitEthernet 0/1/0/1
RP/0/RP0/CPU0:router(config-ospfv3-ar-if)# priority 4
```

Related Commands

| Command | Description |
|------------------------------------------------|-----------------------------------------------------------------------------------------|
| neighbor (OSPFv3), on page 965 | Configures OSPFv3 routers interconnecting to nonbroadcast networks. |
| network (OSPFv3), on page 967 | 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 5.2.1 | 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 **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/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router ospfv3 100
RP/0/RP0/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 | |
|---------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>ipv6-prefix</i> | 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. |
| <i>/ prefix-length</i> | 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 <i>number</i> | (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 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.

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/RP0/CPU0:router(config)# router ospfv3 201  
RP/0/RP0/CPU0:router(config-ospfv3)# area 0  
RP/0/RP0/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. |
|------------|-------------------------------------------|

| | |
|--------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>process-id</i> | <p>For the bgp keyword, an autonomous system number has the following ranges:</p> <ul style="list-style-type: none"> • 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. <p>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.</p> <p>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.</p> |
| metric <i>metric-value</i> | (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 } | <p>(Optional) Specifies the external link type associated with the route advertised into the OSPF routing domain. It can be one of two values:</p> <ul style="list-style-type: none"> • 1—Type 1 external route • 2—Type 2 external route <p>If no metric-type is specified, the default is Type 2 external routes.</p> |
| tag <i>t ag-value</i> | (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 <i>policy-name</i> | (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 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.



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 redistribution 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/RP0/CPU0:router(config)# router ospfv3 109
RP/0/RP0/CPU0:router(config-ospfv3)# redistribute isis level-1
```

| Related Commands | Command | Description |
|------------------|---------------------------------------------------------------------|----------------------------------------------------------------------------|
| | default-information originate (OSPFv3), on page 935 | Redistributes routes into a routing domain. |
| | distribute-list prefix-list out, on page 945 | 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 | <i>seconds</i> 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 | <p>If this command is not specified in interface configuration mode, then the interface adopts the retransmit interval parameter specified by the area.</p> <p>If this command is not specified in area configuration mode, then the interface adopts the retransmit interval parameter specified by the process.</p> <p>If this command is not specified at any level, then the default retransmit interval is 5 seconds.</p> | |
| Command Modes | <p>Interface configuration</p> <p>Area configuration</p> <p>Router OSPFv3 configuration</p> <p>Virtual-link configuration</p> | |
| Command History | Release | Modification |
| | Release 5.0.0 | This command was introduced. |
| Usage Guidelines | <p>To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.</p> <p>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.</p> <p>The setting of this parameter should be conservative, or needless retransmission results. The value should be larger for serial lines and virtual links.</p> | |
| Task ID | Task ID | Operations |
| | ospf | read, write |
| Examples | <p>The following example shows how to set the retransmit interval value to 8 seconds while in interface configuration mode:</p> | |

```
RP/0/RP0/CPU0:router(config)# router ospfv3 1  
RP/0/RP0/CPU0:router(config-ospfv3)# area 0  
RP/0/RP0/CPU0:router(config-ospfv3-ar)# interface GigabitEthernet 0/2/0/0  
RP/0/RP0/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 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.

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/RP0/CPU0:router(config)# router ospfv3 109  
RP/0/RP0/CPU0:router(config-ospfv3)# router-id 10.0.0.10
```

Related Commands

| Command | Description |
|---------------------------------------------------|---------------------------------------------------------------------|
| clear ospfv3 process, on page 921 | 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 XR Config 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 | <i>process-name</i> 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 | XR Config |
|----------------------|-----------|

| Command History | Release | Modification |
|------------------------|----------------|------------------------------|
| | 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. |
|-------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

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/RP0/CPU0:router(config)# router ospfv3 1
```

| Related Commands | Command | Description |
|-------------------------|-------------------------------------------------|--------------------------------------|
| | area (OSPFv3), on page 916 | Defines an OSPFv3 area. |
| | interface (OSPFv3), on page 956 | Defines an OSPFv3 interface by type. |

show ospfv3

To display general information about Open Shortest Path First Version 3 (OSPFv3) routing processes, use the **show ospfv3** command in XR EXEC mode.

```
show ospfv3 [process-name]
```

| | | |
|---------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------|
| Syntax Description | <i>process-name</i> (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 | XR EXEC | |
| Command History | Release | Modification |
| | Release 5.0.0 | This command was introduced. |
| | Release 5.2.1 | The output of the command was modified to include OSPFv3 graceful shutdown status. |
| Usage Guidelines | To use this command, you must be in a user group associated with a task 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/RP0/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 msec
  Minimum hold time between two consecutive SPF's 10000 msec
  Maximum wait time between two consecutive SPF's 10000 msec
  Initial LSA throttle delay 0 msec
  Minimum hold time for LSA throttle 5000 msec
  Maximum wait time for LSA throttle 5000 msec
  Minimum LSA arrival 1000 msec
  LSA group pacing timer 240 secs
  Interface flood pacing timer 33 msec
  Retransmission pacing timer 66 msec
```

```

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

```

The following is a sample output from the **show ospfv3** command when OSPFv3 graceful shutdown has been initiated but not yet completed:

```

RP/0/RP0/CPU0:router# show ospfv3 1

Routing Process "ospfv3 test" with ID 3.3.3.3 transitioning to OSPFV3_ADMIN_DOWN state

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 msec
Minimum hold time between two consecutive SPFs 10000 msec
Maximum wait time between two consecutive SPFs 10000 msec
Initial LSA throttle delay 0 msec
Minimum hold time for LSA throttle 5000 msec
Maximum wait time for LSA throttle 5000 msec
Minimum LSA arrival 1000 msec
LSA group pacing timer 240 secs
Interface flood pacing timer 33 msec
Retransmission pacing timer 66 msec
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

```

The following is a sample output from the **show ospfv3** command when OSPFv3 graceful shutdown is completed:

```

RP/0/RP0/CPU0:router# show ospfv3 1

Routing Process "ospfv3 test" with ID 3.3.3.3 in OSPFV3_ADMIN_DOWN state
%ROUTING-OSPFv3-6-GRACEFUL_SHUTDOWN : Shutdown Complete

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 msec
Minimum hold time between two consecutive SPF's 10000 msec
Maximum wait time between two consecutive SPF's 10000 msec
Initial LSA throttle delay 0 msec
Minimum hold time for LSA throttle 5000 msec
Maximum wait time for LSA throttle 5000 msec
Minimum LSA arrival 1000 msec
LSA group pacing timer 240 secs
Interface flood pacing timer 33 msec
Retransmission pacing timer 66 msec
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 95: 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 SPF's | Minimum hold time between consecutive SPF's. |
| Maximum wait time between two consecutive SPF's | Maximum wait time between consecutive SPF's. |
| 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. |

| Field | Description |
|-----------------------------------------|--------------------------------------------|
| 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. |
| Number of interfaces in this area is | Number of interfaces in the area. |
| SPF algorithm executed <i>n</i> 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. |

This is sample output from the **show ospfv3** command to verify that (Non-stop routing (NSR) is enabled:

```
RP/0/RP0/CPU0:router#show ospfv3
```

```
Routing Process "ospfv3 100" with ID 3.3.3.3
NSR (Non-stop routing) is Enabled
It is an area border and autonomous system boundary router
Redistributing External Routes from,
  bgp 100
  Maximum number of redistributed prefixes 10240
  Threshold for warning message 75%
Initial SPF schedule delay 5000 msec
Minimum hold time between two consecutive SPF's 10000 msec
Maximum wait time between two consecutive SPF's 10000 msec
Initial LSA throttle delay 0 msec
Minimum hold time for LSA throttle 5000 msec
Maximum wait time for LSA throttle 5000 msec
Minimum LSA arrival 1000 msec
LSA group pacing timer 240 secs
Interface flood pacing timer 33 msec
Retransmission pacing timer 66 msec
Maximum number of configured interfaces 512
Maximum number of configured paths 16
Number of external LSA 0. Checksum Sum 00000000
Number of areas in this router is 15. 15 normal 0 stub 0 nssa
Auto cost is enabled. Reference bandwidth 100
```

Related Commands

| Command | Description |
|--------------------------------------------|---------------------------------------|
| router ospfv3, on page 990 | 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 XR EXEC mode.

```
show ospfv3 [process-name] border-routers [router-id]
```

| Syntax Description | <p><i>process-name</i> (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.</p> <p><i>router-id</i> (Optional) 32-bit router ID value specified in four-part, dotted-decimal notation.</p> | | | | |
|---------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------|--------------|---------------|------------------------------|
| Command Default | No default behavior or values | | | | |
| Command Modes | XR EXEC | | | | |
| Command History | <table border="1"> <thead> <tr> <th>Release</th> <th>Modification</th> </tr> </thead> <tbody> <tr> <td>Release 5.0.0</td> <td>This command was introduced.</td> </tr> </tbody> </table> | Release | Modification | Release 5.0.0 | This command was introduced. |
| Release | Modification | | | | |
| 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 | <table border="1"> <thead> <tr> <th>Task ID</th> <th>Operations</th> </tr> </thead> <tbody> <tr> <td>ospf</td> <td>read</td> </tr> </tbody> </table> | Task ID | Operations | ospf | read |
| Task ID | Operations | | | | |
| ospf | read | | | | |

Examples

The following is sample output from the **show ospfv3 border-routers** command:

```
RP/0/RP0/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 96: 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 990 | 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 XR 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

| | |
|----------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>process-name</i> | (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. |
| <i>area-id</i> | (Optional) Area number used to define the particular area. |
| adv-router [<i>router-id</i>] | (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). |
| <i>link-state-id</i> | (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

XR EXEC

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | 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.

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/RP0/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 97: 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. |
| 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/RP0/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 98: 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/RP0/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 99: 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. 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/RP0/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 100: 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. 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/RP0/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 101: 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.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. 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/RP0/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 102: 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.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/RP0/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 103: 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. 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/RP0/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 104: 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. 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: <ul style="list-style-type: none">• 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 990 | 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 XR EXEC mode.

```
show ospfv3 [process-name] [area-id] flood-list [type interface-path-id]
```

| Syntax Description | |
|--------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>process-name</i> | (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. |
| <i>area-id</i> | (Optional) Area number used to define the particular area. |
| <i>type</i> | Interface type. For more information, use the question mark (?) online help function. |
| <i>interface-path-id</i> | Physical interface or virtual interface. |
| | <p>Note Use the show interfaces command to see a list of all interfaces currently configured on the router.</p> <p>For more information about the syntax for the router, use the question mark (?) online help function.</p> |

| Command Default | |
|-----------------|-------------------------------|
| | No default behavior or values |

| Command Modes | |
|---------------|---------|
| | XR EXEC |

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | 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. |

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/RP0/CPU0:router# show ospfv3 flood-list GigabitEthernet 0/3/0/0
```

```
Flood Lists for OSPFv3 1
```

show ospfv3 flood-list

```
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 105: 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 990 | 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 XR EXEC mode.

```
show ospfv3 [process-name] [area-id] interface [type interface-path-id]
```

Syntax Description

| | |
|--------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>process-name</i> | (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. |
| <i>area-id</i> | (Optional) Area number used to define the particular area. |
| <i>type</i> | Interface type. For more information, use the question mark (?) online help function. |
| <i>interface-path-id</i> | 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

XR EXEC

Command History

| Release | Modification |
|---------------|------------------------------|
| 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.

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/ /0/0 is specified:

```

RP/0/RP0/CPU0:router# show ospfv3 interface GigabitEthernet 0/

/

/0//0/0 is up, line protocol is
Link Local address fe80:::, Interface ID 2
Area 0, Process ID , Instance ID 0, Router ID ...
Network Type BROADCAST, Cost:
BFD enabled, interval 300 msec, multiplier 5
Transmit Delay is 1 sec, State BDR, Priority 1
Designated Router (ID) ..., local address fe80:::
Backup Designated router (ID) ..., local address fe80:::
Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
  Hello due in 00:00:06
Index 0//1, flood queue length 0
Next 0(0)/0(0)/0(0)
Last flood scan length is , maximum is
Last flood scan time is 0 msec, maximum is msec
Neighbor Count is 1, Adjacent neighbor count is 1
  Adjacent with neighbor ... (Designated Router)
Suppress hello for 0 neighbor(s)

```

This table describes the significant fields shown in the display.

Table 106: 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 990 | 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 XR EXEC mode.

show ospfv3 [*process-name*] **message-queue**

This command has no keywords or arguments.

Command Default None

Command Modes XR EXEC

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | 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 |
|---------|---------|-----------|
| | ospf | read |

This is sample output from the **show ospfv3 message-queue** command:

```
RP/0/RP0/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 XR EXEC mode.

show ospfv3 [*process-name*] [*area-id*] **neighbor** [*type interface-path-id*] [*neighbor-id*] [**detail**]

| Syntax Description | |
|--------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>process-name</i> | (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. |
| <i>area-id</i> | (Optional) Area ID. If you do not specify an area, all areas are displayed. |
| <i>type</i> | Interface type. For more information, use the question mark (?) online help function. |
| <i>interface-path-id</i> | 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. |
| <i>neighbor-id</i> | (Optional) Neighbor router ID. |
| detail | (Optional) Displays all neighbors given in detail (lists all neighbors). |

Command Default No default behavior or values

Command Modes XR EXEC

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | 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.

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/RP0/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 107: 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/RP0/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 108: 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/RP0/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 109: 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/RP0/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 110: 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/RP0/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 111: show ospfv3 neighbor GigabitEthernet 0/3/0/1 detail Field Descriptions

| Field | Description |
|-------------|----------------------------------------------------------------|
| Neighbor | Neighbor router ID. |
| In the area | Area and interface through which the OSPFv3 neighbor is known. |

| Field | Description |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------|
| 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 990 | 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 XR EXEC mode.

```
show ospfv3 [process-name] [area-id] request-list [type interface-path-id] [neighbor-id]
```

Syntax Description

| | |
|--------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>process-name</i> | (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. |
| <i>area-id</i> | (Optional) Area ID. If you do not specify an area, all areas are displayed. |
| <i>type</i> | (Optional) Interface type. For more information, use the question mark (?) online help function. |
| <i>interface-path-id</i> | (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. |
| <i>neighbor-id</i> | (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

XR EXEC

Command History

| Release | Modification |
|---------------|------------------------------|
| 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.

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/RP0/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 112: 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 990 | Configures an OSPFv3 routing process. |
| show ospfv3 retransmission-list, on page 1024 | 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 XR EXEC mode.

```
show ospfv3 [process-name ] [area-id] retransmission-list [type interface-path-id] [neighbor-id]
```

Syntax Description

| | |
|--------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>process-name</i> | (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. |
| <i>area-id</i> | (Optional) Area ID. If you do not specify an area, all areas are displayed. |
| <i>type</i> | (Optional) Interface type. For more information, use the question mark (?) online help function. |
| <i>interface-path-id</i> | (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. |
| <i>neighbor-id</i> | (Optional) IP address of the OSPFv3 neighbor. |

Command Default

No default behavior or values

Command Modes

XR EXEC

Command History

| Release | Modification |
|---------------|------------------------------|
| 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.

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/RP0/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 113: 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 990 | Configures an OSPFv3 routing process. |
| show ospfv3 request-list, on page 1022 | 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 XR EXEC mode.

```
show ospfv3 [process-name] routes [{external | connected}] [ipv6-prefix/prefix-length]
show ospfv3 [process-name] routes summary
```

Syntax Description

| | |
|------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>process-name</i> | (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. |
| <i>ipv6-prefix</i> | (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. |
| <i>/ prefix-length</i> | (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

XR EXEC

Command History

| Release | Modification |
|---------------|------------------------------|
| 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.

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/RP0/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 114: 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 990 | Configures an OSPFv3 routing process. |

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 XR EXEC mode.

show ospfv3 [*process-name*] **summary-prefix**

| | |
|---------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax Description | <i>process-name</i> (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 | XR EXEC |
|----------------------|---------|

| Command History | Release | Modification |
|------------------------|----------------|------------------------------|
| | 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.

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/RP0/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 115: 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 990 | Configures an OSPFv3 routing process. |
| summary-prefix (OSPFv3), on page 1038 | 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 XR EXEC mode.

show ospfv3 [*process-name*] **virtual-links**

| | |
|---------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax Description | <i>process-name</i> (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 | XR EXEC |
|----------------------|---------|

| | | |
|------------------------|----------------|------------------------------|
| Command History | Release | Modification |
| | 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. |
|-------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

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/RP0/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 116: 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 990 | 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 XR EXEC mode.

```
show protocols [{afi-all | ipv4 | ipv6}] [{allprotocol}]
```

| 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. |
| <i>protocol</i> | (Optional) Specifies a routing protocol. For the IPv4 address family, the options are: <ul style="list-style-type: none"> • bgp • eigrp • isis • ospf • rip For the IPv6 address family, the options are: <ul style="list-style-type: none"> • bgp • eigrp • isis • ospfv3 |

Command Default The default address family is IPv4.

Command Modes XR EXEC

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | 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 | Operations |
|---------|---------|------------|
| | ospf | read |

Examples The following example is sample output from the **show protocols** command:

```
RP/0/RP0/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 117: 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. |

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 <i>totally stubby</i> areas. | | | | |
|---------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------|--------------|---------------|------------------------------|
| Command Default | No stub area is defined. | | | | |
| Command Modes | Area configuration | | | | |
| Command History | <table border="1"> <thead> <tr> <th>Release</th> <th>Modification</th> </tr> </thead> <tbody> <tr> <td>Release 5.0.0</td> <td>This command was introduced.</td> </tr> </tbody> </table> | Release | Modification | Release 5.0.0 | This command was introduced. |
| Release | Modification | | | | |
| 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.

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/RP0/CPU0:router(config)# router ospfv3 201
RP/0/RP0/CPU0:router(config-ospfv3)# area 5
RP/0/RP0/CPU0:router(config-ospfv3-ar)# stub
```

```
RP/0/RP0/CPU0:router(config-ospfv3-ar)# default-cost 20
```

Related Commands

| Command | Description |
|----------------------------------------------------|-----------------------------------------------------------------------|
| default-cost (OSPFv3), on page 933 | 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. |
| summary-lsa | <p>If enabled, summary LSAs are advertised with modified metric when stub-router is active. This configuration is applicable to max-metric mode.</p> <p>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).</p> <p>If enabled and metric is not explicitly configured, the default metric for summary LSAs when stub-router active is 16711680 (0xFF0000).</p> |

| | |
|---------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| external-lsa | <p>If enabled, external LSAs are advertised with modified metric when stub-router is active. This configuration is applicable to max-metric mode.</p> <p>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).</p> <p>If enabled and metric is not explicitly configured, the default metric for external LSAs when stub-router active is 16711680 (0xFF0000).</p> |
| include-stub | <p>If enabled, intra-area-prefix LSAs that are referencing router LSA are advertised with maximum metric (0xffff) when stub-router is active.</p> <p>Intra-area-prefix LSAs that are referencing network LSA do not change metric</p> <p>Can be used in r-bit and max-metric modes.</p> <p>/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.</p> |

Command Default

Disabled.

Command Modes

Router OSPFv3 configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| 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.

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 |

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 | |
|------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>ipv6-prefix</i> | 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. |
| <i>/ prefix-length</i> | 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 OSPFv3 protocol.

Command Modes Router OSPFv3 configuration

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | 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.

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/RP0/CPU0:router(config-ospfv3)# summary-prefix 4004:f000::/32
```

Related Commands

| Command | Description |
|---------------------------------------------|---------------------------------------------------------|
| range (OSPFv3), on page 980 | 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 | <i>milliseconds</i> 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 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. |
|-------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

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/RP0/CPU0:router(config)# router ospfv3 1
RP/0/RP0/CPU0:router(config-ospfv3)# timers throttle lsa all 200 10000 45000
RP/0/RP0/CPU0:router(config-ospfv3)# timers lsa arrival 2000
```

| Related Commands | Command | Description |
|-------------------------|----------------------------------------------------------------|-----------------------------------------------------|
| | timers throttle lsa all (OSPFv3), on page 1047 | 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 | <i>milliseconds</i> Time (in milliseconds) at which LSAs in the flooding queue are paced in between updates. Range is 5 milliseconds to 100 milliseconds. | | | | |
|---------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------|--------------|---------------|------------------------------|
| Command Default | <i>milliseconds: 33</i> | | | | |
| Command Modes | Router OSPFv3 configuration | | | | |
| Command History | <table border="1"> <thead> <tr> <th>Release</th> <th>Modification</th> </tr> </thead> <tbody> <tr> <td>Release 5.0.0</td> <td>This command was introduced.</td> </tr> </tbody> </table> | Release | Modification | Release 5.0.0 | This command was introduced. |
| Release | Modification | | | | |
| 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.

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/RP0/CPU0:router(config)# router ospfv3 1
RP/0/RP0/CPU0:router(config-ospfv3)# timers pacing flood 55
```

Related Commands

| Command | Description |
|-------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------|
| show ospfv3, on page 991 | Displays general information about OSPFv3 routing processes. |
| timers pacing lsa-group, on page 1043 | 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 1045 | 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 | <i>seconds</i> Interval (in seconds) at which LSAs are grouped and refreshed, checksummed, or aged. Range is 10 to 1800 seconds. | |
| Command Default | <i>seconds</i> : 240 OSPFv3 LSA group pacing is enabled by default. | |
| Command Modes | Router OSPFv3 configuration | |
| Command History | Release | Modification |
| | Release 5.0.0 | This command was introduced. |
| Usage Guidelines | <p>To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.</p> <p>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.</p> <p>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).</p> <p>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.</p> | |
| 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/RP0/CPU0:router(config)# router ospfv3 1
RP/0/RP0/CPU0:router(config-ospfv3)# timers pacing lsa-group 60
```

Related Commands

| Command | Description |
|-------------------------------------------------------------|--------------------------------------------------------------|
| show ospfv3, on page 991 | Displays general information about OSPFv3 routing processes. |
| timers pacing flood, on page 1041 | Configures LSA flood packet pacing. |
| timers pacing retransmission , on page 1045 | 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 | <i>milliseconds</i> Time (in milliseconds) at which LSAs in the retransmission queue are paced. Range is 5 milliseconds to 100 milliseconds. | | | | |
|---------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------|--------------|---------------|------------------------------|
| Command Default | <i>milliseconds</i> : 66 | | | | |
| Command Modes | Router OSPFv3 configuration | | | | |
| Command History | <table border="1"> <thead> <tr> <th>Release</th> <th>Modification</th> </tr> </thead> <tbody> <tr> <td>Release 5.0.0</td> <td>This command was introduced.</td> </tr> </tbody> </table> | Release | Modification | Release 5.0.0 | This command was introduced. |
| Release | Modification | | | | |
| 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.

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/RP0/CPU0:router(config)# router ospfv3 1
```

```
RP/0/RP0/CPU0:router(config-ospfv3)# timers pacing retransmission 55
```

Related Commands

| Command | Description |
|-------------------------------------------------------|-----------------------------------------------------------------------------------------------------------|
| show ospfv3, on page 991 | Displays general information about OSPFv3 routing processes. |
| timers pacing flood, on page 1041 | Configures LSA flood packet pacing. |
| timers pacing lsa-group, on page 1043 | 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 | |
|-----------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>start-interval</i> | 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. |
| <i>hold-interval</i> | Incremental time (in milliseconds). This value is used to calculate the subsequent rate limiting times for LSA generation. Range is 1 to 600000 milliseconds. |
| <i>max-interval</i> | Maximum wait time (in milliseconds) between generation of the same LSA. Range is 1 to 600000 milliseconds. |

| Command Default | |
|-----------------|------------------------------------------|
| | <i>start-interval</i> : 500 milliseconds |
| | <i>hold-interval</i> : 5000 milliseconds |
| | <i>max-interval</i> : 5000 milliseconds |

| Command Modes | |
|---------------|-----------------------------|
| | Router OSPFv3 configuration |

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | 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. |

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/RP0/CPU0:router(config)# router ospfv3 1  
RP/0/RP0/CPU0:router(config-ospfv3)# timers throttle lsa all 200 10000 45000  
RP/0/RP0/CPU0:router(config-ospfv3)# timers lsa arrival 2000
```

Related Commands

| Command | Description |
|--------------------------------------------------|---------------------------------------------------------------------------------------------|
| show ospfv3, on page 991 | Displays general information about OSPFv3 routing processes. |
| timers lsa arrival, on page 1040 | 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 | | |
|--------------------|---------------------|------------------------------------------------------------------------------------------------------------------|
| | <i>spf-start</i> | Initial SPF schedule delay (in milliseconds). Range is 1 to 600000 milliseconds. |
| | <i>spf-hold</i> | Minimum hold time (in milliseconds) between two consecutive SPF calculations. Range is 1 to 600000 milliseconds. |
| | <i>spf-max-wait</i> | Maximum wait time (in milliseconds) between two consecutive SPF calculations. Range is 1 to 600000 milliseconds. |

| Command Default | |
|-----------------|------------------------------------------|
| | <i>spf-start</i> : 5000 milliseconds |
| | <i>spf-hold</i> : 10000 milliseconds |
| | <i>spf-max-wait</i> : 10000 milliseconds |

| Command Modes | |
|---------------|-----------------------------|
| | Router OSPFv3 configuration |

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | 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.

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/RP0/CPU0:router(config)# router ospfv3 1  
RP/0/RP0/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 | Parameter | Description |
|--------------------|-------------|------------------------------------------------------------------------------------------------------------------------------------------------------|
| | size | Deletes existing buffer and creates one with <i>N</i> entries. |
| | buffer_name | Specifies a buffer from one of the 15 listed buffers. Refer Table 118: Buffer Types, on page 1051 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 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.

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 118: Buffer Types

| Name | Description |
|-----------|-----------------------|
| adj | adjacency |
| adj_cycle | dbd/flood events/pkts |
| config | config events |
| errors | errors |
| events | mda/rtrid/bfd |
| 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 Operations**

| Task ID | Operations |
|---------|----------------|
| ospf | read, write |

Examples

This example shows how to set 1024 error trace entries:

```
RP/0/RP0/CPU0:router#configure
RP/0/RP0/CPU0:router(config)#router ospfv3 osp3
RP/0/RP0/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/RP0/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 | <i>seconds</i> 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 5.0.0 | This command was introduced. |
| Usage Guidelines | <p>To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.</p> <p>Link-state advertisements (LSAs) in the update packet must have their ages incremented by the amount specified in the <i>seconds</i> argument before transmission. The value assigned should take into account the transmission and propagation delays for the interface.</p> <p>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.</p> | |
| Task ID | Task ID | Operations |
| | ospf | read, write |
| Examples | <p>The following example shows how to configure a transmit delay of 3 seconds for GigabitEthernet interface 0/3/0/0:</p> <pre>RP/0/RP0/CPU0:router(config)# router ospfv3 1 RP/0/RP0/CPU0:router(config-ospfv3)# area 0 RP/0/RP0/CPU0:router(config-ospfv3-ar)# interface GigabitEthernet 0/3/0/0 RP/0/RP0/CPU0:router(config-ospfv3-ar-if)# transmit-delay 3</pre> | |

Related Commands

| Command | Description |
|------------------------------------------|------------------------------------------------------------|
| show ospfv3, on page 991 | 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 | <i>router-id</i> 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 | <table border="1"> <thead> <tr> <th>Release</th> <th>Modification</th> </tr> </thead> <tbody> <tr> <td>Release 5.0.0</td> <td>This command was introduced.</td> </tr> </tbody> </table> | Release | Modification | Release 5.0.0 | This command was introduced. |
| Release | Modification | | | | |
| Release 5.0.0 | This command was introduced. | | | | |
| Usage Guidelines | <p>To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.</p> <p>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.</p> <p>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.</p> <p>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.</p> <p>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.</p> | | | | |
| Task ID | <table border="1"> <thead> <tr> <th>Task ID</th> <th>Operations</th> </tr> </thead> <tbody> <tr> <td>ospf</td> <td>read, write</td> </tr> </tbody> </table> | Task ID | Operations | ospf | read, write |
| Task ID | Operations | | | | |
| ospf | read, write | | | | |
| Examples | <p>The following example shows how to establish a virtual link with default values for all optional parameters:</p> <pre>RP/0/RP0/CPU0:router(config)# router ospfv3 201</pre> | | | | |

```
RP/0/RP0/CPU0:router(config-ospfv3)# area 1
RP/0/RP0/CPU0:router(config-ospfv3-ar)# virtual-link 10.3.4.5
```

Related Commands

| Command | Description |
|-------------------------------------------------------|---------------------------------------------------------------------------------------|
| show ospfv3, on page 991 | Displays general information about OSPF routing processes. |
| transmit-delay (OSPFv3), on page 1053 | Sets the estimated time required to send a link-state update packet on the interface. |



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 Softwaremodule* in *Routing Configuration Guide for Cisco NCS 6000 Series Routers*.

- [address-family next-hop dampening disable](#), on page 1059
- [clear route](#), on page 1060
- [lcc](#), on page 1062
- [rcc](#), on page 1063
- [recursion-depth-max](#), on page 1065
- [router rib](#), on page 1066
- [rump always-replicate](#), on page 1067
- [show lcc statistics](#), on page 1068
- [show rcc](#), on page 1070
- [show rcc statistics](#), on page 1072
- [show rib](#), on page 1074
- [show rib afi-all](#), on page 1076
- [show rib attributes](#), on page 1078
- [show rib client-id](#), on page 1079
- [show rib clients](#), on page 1081
- [show rib extcomms](#), on page 1082
- [show rib firsthop](#), on page 1084
- [show rib history](#), on page 1086
- [show rib next-hop](#), on page 1088
- [show rib opaques](#), on page 1090
- [show rib protocols](#), on page 1092
- [show rib recursion-depth-max](#), on page 1094
- [show rib statistics](#), on page 1096
- [show rib tables](#), on page 1098
- [show rib trace](#), on page 1099
- [show route](#), on page 1101
- [show route backup](#), on page 1106
- [show route best-local](#), on page 1109
- [show route connected](#), on page 1111

- [show route local](#), on page 1113
- [show route longer-prefixes](#), on page 1115
- [show route next-hop](#), on page 1117
- [show route quarantined](#), on page 1119
- [show route resolving-next-hop](#), on page 1121
- [show route static](#), on page 1123
- [show route summary](#), on page 1125

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 | <p>ipv4 Specifies IP Version 4 (IPv4) address prefixes.</p> <p>ipv6 Specifies IP Version 6 (IPv6) address prefixes.</p> | |
| Command Default | RIB next-hop dampening is enabled. | |
| Command Modes | Router configuration | |
| Command History | Release | Modification |
| | 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 | Operations |
| | rib | read, write |

Examples

The following example shows how to disable RIB next-hop dampening for IPv6 address families:

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router rib
RP/0/RP0/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 XR EXEC mode.

```
clear route {ipv4 | ipv6} {unicast | safi-all} [topology topo-name] [ip-address mask]
```

Syntax Description

| | |
|----------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| ipv4 | Specifies IP Version 4 address prefixes. |
| ipv6 | Specifies IP Version 6 address prefixes. |
| | Specifies IP Version 4 and IP Version 6 address prefixes. |
| unicast | Specifies unicast address prefixes. |
| safi-all | Specifies unicast address prefixes. |
| topology <i>topo-name</i> | (Optional) Specifies topology table information and name of the topology table. |
| <i>ip-address node-id</i> | (Optional) Clears hardware resource counters from the designated node. The <i>node-id</i> argument is entered in the <i>rack/slot/module</i> notation. |
| <i>ip-address</i> | Network IP address about which routing information should be displayed. |
| <i>mask</i> | 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 Modes

XR EXEC

Command History

| Release | Modification |
|---------------|------------------------------|
| 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.

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/RP0/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/RP0/CPU0:router# clear route ipv4 unicast
```

Related Commands

| Command | Description |
|------------------------------------------|--------------------------------------------------|
| show route, on page 1101 | Displays the current state of the routing table. |

lcc

To enable Label Consistency Checker (lcc) background scan for IPv6 or IPv4 labels, use the **lcc enable** command in XR Config 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 | Parameter | Description |
|--------------------|-----------------------------------|------------------------------------------------------------------------------------------|
| | ipv4 | Specifies IP Version 4 address prefixes. |
| | ipv6 | Specifies IP Version 6 address prefixes. |
| | unicast | Specifies unicast address prefixes. |
| | period <i>milliseconds</i> | Specifies the period between scans in milliseconds. Range is 100 to 600000 milliseconds. |

Command Default Label consistency checker is disabled.

Command Modes XR Config

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | 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 |
|---------|---------|----------------|
| | ipv4 | read, write |
| | ipv6 | read, write |

This example shows how to enable lcc for IPv6 labels:

```
RP/0/RP0/CPU0:router#configure
RP/0/RP0/CPU0:router(config)#lcc ipv6 unicast enable
```

rcc

To enable Route Consistency Checker (rcc) background scan for IPv6 or IPv4 routes, use the **rcc enable** command in XR Config mode. To disable rcc background scan, use the **no** form of this command.

```
rcc {ipv4 | ipv6} unicast {enable | period milliseconds}
no rcc {ipv4 | ipv6} unicast {enable | period milliseconds}
```

| Syntax Description | Parameter | Description |
|--------------------|-----------------------------------|------------------------------------------------------------------------------------------|
| | ipv4 | Specifies IP Version 4 address prefixes. |
| | ipv6 | Specifies IP Version 6 address prefixes. |
| | unicast | Specifies unicast address prefixes. |
| | period <i>milliseconds</i> | Specifies the period between scans in milliseconds. Range is 100 to 600000 milliseconds. |

Command Default Route consistency checker is disabled.

Command Modes XR Config

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | 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.

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:

```
RP/0/RP0/CPU0:router#configure
RP/0/RP0/CPU0:router(config)#rcc ipv6 unicast enable
```

This example shows how to enable rcc with a scan period of 500 milliseconds for IPv6 unicast:

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

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/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router rib
RP/0/RP0/CPU0:router(config-rib)# recursion-depth-max 12
```

router rib

To enter Routing Information Base (RIB) configuration mode, use the **router rib** command in XR Config 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 XR Config

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | 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 | 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/RP0/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 | <i>access-list-name</i> (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 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. |
|-------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

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/RP0/CPU0:router(config)# router rib  
RP/0/RP0/CPU0:router(config-rib)# address-family ipv4  
RP/0/RP0/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 XR 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 | <i>scan-id-value</i> | 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 XR EXEC

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | 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 |
|---------|---------|-----------|
| | ipv4 | read |
| | ipv6 | read |

This example shows background scan statistics for AFI-SAFI mpls6-unicast:

```
RP/0/RP0/CPU0:router#show lcc ipv6 unicast statistics

Background Scan Statistics for AFI-SAFI mpls6-unicast:
=====

Scan enabled:                False
Current scan-id:             0
Configured period:           60
Scan triggered:              False
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 mpls v6-unicast:

```
=====
```

End Of Logs

This example shows background scan statistics for AFI-SAFI mpls v4-unicast:

```
RP/0/RP0/CPU0:router#show lcc ipv4 unicast statistics
```

```
Background Scan Statistics for AFI-SAFI mpls v4-unicast:
```

```
=====
```

```

Scan enabled:           False
Current scan-id:       0
Configured period:    60
Scan triggered:        False
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 mpls v4-unicast:

```
=====
```

End Of Logs

show rcc

To display route consistency checker (RCC) information, use the **show rcc** command in XR EXEC mode.

```
show rcc {ipv4 | ipv6} unicast [{prefix netmask }]
```

| 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. |

Command Default No default behavior or values

Command Modes XR EXEC

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | 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 | Operations |
|---------|---------|------------|
| | ipv4 | read |

Examples

The following is sample output from the **show rcc** command:

```
RP/0/RP0/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 XR 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 | <i>scan-id-value</i> | 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 XR EXEC

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | 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 |
|---------|---------|-----------|
| | ipv4 | read |
| | ipv6 | read |

This example shows background scan statistics for AFI-SAFI IPv6 unicast:

```
RP/0/RP0/CPU0:router#show rcc ipv6 unicast statistics

Background Scan Statistics for AFI-SAFI ipv6-unicast:
=====

Scan enabled:                False
Current scan-id:             0
Configured period:           60
Scan triggered:              False
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/RP0/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 rib

To display Routing Information Base (RIB) data, use the **show rib** command in XR EXEC mode.

```
show rib {ipv4 | ipv6} {unicast} [{firsthop | [type interface-path-id]} | next-hop | [type interface-path-id]} | opaques | {attribute | ip-nexthop | ipfrr | summary} | 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. |
| 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. |
| ipfrr | (Optional) Specifies IP fast reroute (IPFRR) opaque data installed in the RIB. |
| summary | (Optional) Specifies a summary of opaque data installed in the RIB. |
| protocols | (Optional) Specifies registered protocols. |
| statistics <i>name</i> | (Optional) Specifies RIB statistics of a given name. |
| standby | (Optional) Specifies standby information. |
| topology <i>topo-name</i> | (Optional) Specifies topology table information and name of the topology table. |
| all | (Optional) Specifies that all topology table information should be displayed. |

Command Default

No default behavior or values

Command Modes

XR EXEC

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | 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 | Operations |
|---------|---------|------------|
| | ipv4 | read |

Examples

| Related Commands | Command | Description |
|------------------|------------------------------------------------|----------------------------------------------|
| | show rib afi-all, on page 1076 | 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 XR EXEC mode.

```
show rib afi-all [attributes] [client-id] [clients] [extcomms] [firsthop] [history] [next-hop]
[opaques] [protocols] [recursion-depth-max] [safi-all] [statistics] [tables] [trace] [unicast]
```

| 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. |
| | next-hop | (Optional) Displays registered next-hop notification addresses. |
| | opaques | (Optional) Displays opaque data installed in RIB. |
| | protocols | (Optional) Displays registered protocols. |
| | recursion-depth-max | (Optional) Displays maximum recursion depth in RIB. |
| | safi-all | (Optional) Displays unicast 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. |

Command Default No default behavior or values

Command Modes XR EXEC

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | 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 | Operations |
|---------|---------|------------|
| | ipv4 | read |

Examples

The following example illustrates the **show rib afi-all attributes** command:

```
RP/0/RP0/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 1074 | 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 XR 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 XR EXEC

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | 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 | Operations |
|---------|---------|------------|
| | rib | read |

Examples

The following is sample output from the **show rib attributes** command:

```
RP/0/RP0/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 XR EXEC mode.

```
show rib client-id id redistribution history [standby]
```

| Syntax Description | |
|------------------------|----------------------------------------------------------------------|
| <i>id</i> | 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 XR EXEC

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | 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 | Operations |
|---------|---------|------------|
| | rib | read |

Examples

The following is sample output from the **show rib client-id** command:

```
RP/0/RP0/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 119: 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 1081 | Displays RIB clients. |

show rib clients

To display Routing Information Base (RIB) clients, use the **show rib clients** command in XR EXEC mode.

```
show rib [{afi-all | ipv4 | ipv6}] clients [{protocols | redistribution [history]}] [standby]
```

| Syntax Description | Parameter | 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 XR EXEC

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | 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.

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 |

show rib extcomms

To display all extended communities installed in the Routing Information Base (RIB), use the **show rib extcomms** command in XR 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 XR EXEC

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | 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 | Operations |
|---------|---------|------------|
| | rib | read |

Examples

The following is sample output from the **show rib extcomms** command:

```
RP/0/RP0/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 120: 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 XR EXEC mode.

```
show rib [{afi-all | ipv4 | ipv6}] [{unicast | safi-all}] firsthop [client-name] [{type interface-path-id | ip-address /prefix-length | ip-address mask | resolved | unresolved | damped}] [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. |
| unicast | (Optional) Specifies unicast address prefixes. This is the default. |
| safi-all | (Optional) Specifies unicast address prefixes. |
| <i>client-name</i> | (Optional) Name of the RIB client. |
| <i>type</i> | Interface type. For more information, use the question mark (?) online help function. |
| <i>interface-path-id</i> | Physical interface or virtual interface. |
| | <p>Note Use the show interfaces command to see a list of all interfaces currently configured on the router.</p> <p>For more information about the syntax for the router, use the question mark (?) online help function.</p> |
| <i>ip-address</i> | (Optional) Network that BGP advertises. |
| <i>/prefix-length</i> | (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. |
| <i>ip-address mask</i> | (Optional) Network mask applied to the <i>ip-address</i> 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 Modes

XR EXEC

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | 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.

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/RP0/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 XR EXEC mode.

```
show rib [{afi-all | ipv4 | ipv6}] history [client-id client-id] [standby]
```

| Syntax Description | Parameter | 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 <i>client-id</i> | (Optional) Specifies the ID of the client. Range for <i>client-id</i> argument is 0 to 4294967295. |
| | standby | (Optional) Displays standby information. |

Command Default No default behavior or values

Command Modes XR EXEC

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | 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.

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/RP0/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 121: 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 XR EXEC mode.

```
show rib [{afi-all | ipv4 | ipv6}] [{unicast | safi-all}] next-hop [client-name] [{type interface-path-id | ip-address /prefix-length | ip-address mask | resolved | unresolved | damped}] [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. |
| unicast | (Optional) Specifies unicast address prefixes. This is the default. |
| safi-all | (Optional) Specifies unicast address prefixes. |
| <i>client-name</i> | (Optional) Name of the RIB client. |
| <i>type</i> | Interface type. For more information, use the question mark (?) online help function. |
| <i>interface-path-id</i> | Physical interface or virtual interface. |
| | <p>Note Use the show interfaces command to see a list of all interfaces currently configured on the router.</p> <p>For more information about the syntax for the router, use the question mark (?) online help function.</p> |
| <i>ip-address</i> | (Optional) Network IP address about which routing information should be displayed. |
| <i>mask</i> | (Optional) Network mask specified in either of two ways: <ul style="list-style-type: none"> • 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. |
| <i>/ prefix-length</i> | (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 XR EXEC

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | 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.

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/RP0/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 XR EXEC mode.

```
show rib [{afi-all | ipv4 | ipv6}] [{unicast | safi-all}] opaques {attribute | ip-nexthop | ipfrr |
summary} [rib-client-name] [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. |
| unicast | (Optional) | Specifies unicast address prefixes. This is the default. |
| safi-all | (Optional) | Specifies unicast 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. |
| summary | | Displays a summary of opaque data installed in the RIB. |
| <i>rib-client-name</i> | (Optional) | Name of the RIB client. |
| standby | (Optional) | Displays standby information. |

Command Default No default behavior or values

Command Modes XR EXEC

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | 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 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 |

Related Commands

| Command | Description |
|-----------------------------------------------|----------------------------------------------------------------------------|
| show route, on page 1101 | Displays current routes information in the Routing Information Base (RIB). |
| show ospf routes, on page 874 | 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 XR EXEC mode.

```
show rib [{afi-all | ipv4 | ipv6}] [{unicast | safi-all}] protocols [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. |
| unicast | (Optional) Specifies unicast address prefixes. This is the default. |
| safi-all | (Optional) Specifies unicast address prefixes. |
| standby | (Optional) Displays standby information. |

Command Modes XR EXEC

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | 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 | Operations |
|---------|---------|------------|
| | rib | read |

Examples The following is sample output from the **show rib protocols** command:

```
RP/0/RP0/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 122: 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 XR 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 XR EXEC

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | 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.

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/RP0/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 123: 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 XR EXEC mode.

```
show rib [{afi-all | ipv4 | ipv6}] [{unicast | safi-all}] statistics [client-name] [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. |
| unicast | (Optional) Specifies unicast address prefixes. This is the default. |
| safi-all | (Optional) Specifies unicast address prefixes. |
| <i>client-name</i> | (Optional) Name of the RIB client. |
| standby | (Optional) Displays standby information. |

Command Modes

XR EXEC

Command History

| Release | Modification |
|---------------|------------------------------|
| 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.

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/RP0/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_SUCCESS 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 124: 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 <i>n</i> light weight messages | Number of lightweight API messages sent from RIB clients. |
| Received <i>n</i> 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 XR 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

XR EXEC

Command History

| Release | Modification |
|---------------|------------------------------|
| 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.

Use the **show rib tables** command to display all tables known to the RIB, including table attributes.

Task ID

| Task ID | Operations |
|---------|------------|
| rib | read |

show rib trace

To display all Routing Information Base (RIB) library call tracer (ltrace) entries, use the **show rib trace** command in XR 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 <i>node-id</i> argument is entered in the <i>rack/slot/module</i> notation. |
| location { all node-id } | (Optional) Displays ltrace entries for the designated node. The <i>node-id</i> argument is entered in the <i>rack/slot/module</i> notation. The all keyword displays ltrace entries for all nodes. |
| Command Default | No default behavior or values |

Command Modes XR EXEC

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | 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 | Operations |
|---------|---------|------------|
| | rib | read |

Examples

The following is sample output from the **show rib trace** command

```
RP/0/RP0/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//CPU0 t1 Create: Management thread
Mar 16 14:59:27.959 rib/ipv4_rib/rib-startup 0//CPU0 t2 Create: Management event manager

Mar 16 14:59:28.346 rib/ipv4_rib/rib-io 0//CPU0 t1 Initialise: RIB server
Mar 16 14:59:28.346 rib/ipv4_rib/rib-io 0//CPU0 t1 Initialise: Client collection
Mar 16 14:59:28.676 rib/ipv4_rib/rib-io 0//CPU0 t1 Initialise: DB collection
Mar 16 14:59:28.693 rib/ipv4_rib/rib-io 0//CPU0 t1 Initialise: Timer tree
Mar 16 14:59:28.694 rib/ipv4_rib/rib-io 0//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//CPU0 t2 Initialise: Debugging routine
Mar 16 14:59:29.128 rib/ipv4_rib/rib-io 0//CPU0 t1 Register: read, select cb functions
Mar 16 14:59:29.137 rib/ipv4_rib/rib-startup 0//CPU0 t1 Register: cerrno DLL name
lib_rib_error.dll
.
.
.
```


show route

To display the current routes in the Routing Information Base (RIB), use the **show route** command in XR EXEC mode.

```
show route [{afi-all | ipv4 | ipv6}] [{unicast | safi-all}] [{protocol [instance] | ip-address [mask] | ip-address/prefix-length}] [standby] [detail]
```

| 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. |
| unicast | (Optional) Specifies unicast address prefixes. This is the default. |
| safi-all | (Optional) Specifies unicast address prefixes. |
| <i>protocol</i> | (Optional) Name of a routing protocol. If you specify a routing protocol, use one of the following keywords: <ul style="list-style-type: none"> • bgp • eigrp • isis • ospf • rip • static • local • connected |
| instance | (Optional) Number or name used to identify an instance of the specified protocol. |
| <i>ip-address</i> | (Optional) Network IP address about which routing information should be displayed. |
| <i>mask</i> | (Optional) Network mask specified in either of two ways: <ul style="list-style-type: none"> • 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. |
| <i>/prefix-length</i> | (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 Modes XR EXEC

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | 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.

When the **afi-all** keyword is used, the *ip-address* and *mask* arguments are not available.

| 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/RP0/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 125: 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). |

| Field | Description |
|----------------------------|-------------------------------------------------------------------------------------------------------------------------------------------|
| [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/RP0/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/RP0/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 126: 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 |

This example is a sample **show route detail** command output that displays path ID and backup-path ID information:

```

RP/0/RP0/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 1125 | Displays the current contents of the routing table in summary format. |
| show rib opaques, on page 1090 | Displays opaque data installed in the Routing Information Base (RIB). |
| show ospf routes, on page 874 | 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 XR EXEC mode.

```
show route [{afi-all | ipv4 | ipv6}] [{unicast | {topology topo-name} | safi-all}] backup [{ip-address
[mask] ip-address /prefix-length }][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. |
| unicast | (Optional) Specifies unicast address prefixes. |
| safi-all safi-all | (Optional) Specifies unicast address prefixes. |
| <i>ip-address</i> | (Optional) Network IP address about which backup routing information should be displayed. |
| <i>mask</i> | (Optional) Network mask specified in either of two ways: <ul style="list-style-type: none"> • 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. |
| <i>/prefix-length</i> | (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 Modes

XR EXEC

Command History

| Release | Modification |
|---------------|------------------------------|
| 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.

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.

| Task ID | Task ID | Operations |
|---------|---------|------------|
| | rib | read |

Examples

The following is sample output from the **show route backup** command:

```
RP/0/RP0/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 127: 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 1101 | 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 XR EXEC mode.

```
show route [{ipv4 | ipv6}] [{unicast | {topology topo-name} | safi-all}] best-local ip-address [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. |
| safi-all | (Optional) | Specifies unicast address prefixes. |
| <i>ip-address</i> | | IP address about which best local information should be displayed. |
| standby | (Optional) | Displays standby information. |

Command Modes XR EXEC

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | 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.

Use the **show route best-local** command to display information about the best local routes in the routing table.

| Task ID | Task ID | Operations |
|---------|---------|------------|
| | rib | read |

Examples

The following is sample output from the **show route best-local** command:

```
RP/0/RP0/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 128: 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 1113 | 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 XR EXEC mode.

```
show route [{afi-all | ipv4 | ipv6}] [{unicast | {topology topo-name} | safi-all}] connected [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. |
| unicast | (Optional) Specifies unicast address prefixes. |
| safi-all | (Optional) Specifies unicast address prefixes. |
| standby | (Optional) Displays standby information. |

| Command Modes | |
|---------------|---------|
| | XR EXEC |

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | 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.

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/RP0/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 129: 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 1125 | 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 XR EXEC mode.

```
show route [{afi-all | ipv4 | ipv6}] [{unicast | {topology topo-name} | safi-all}] local [{type interface
-path-id}] [ 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. |
| unicast | (Optional) Specifies unicast address prefixes. |
| safi-all | (Optional) Specifies unicast address prefixes. |
| <i>type</i> | Interface type. For more information, use the question mark (?) online help function. |
| <i>interface-path-id</i> | 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 Modes XR EXEC

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | 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.

Use the **show route local** command to display information about local routes in the routing table.

| Task ID | Task ID | Operations |
|---------|---------|------------|
| | rib | read |

Examples The following is sample output from the **show route local** command:

```
RP/0/RP0/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 130: 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 1111 | 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 XR EXEC mode.

```
show route [{ipv4 | ipv6}] [{unicast | {topology topo-name} | safi-all}] longer-prefixes {ip-address
mask ip-address/prefix-length} [ 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. |
| safi-all | (Optional) Specifies unicast address prefixes. |
| <i>ip-address</i> | Network IP address about which routing information should be displayed. |
| <i>mask</i> | Network mask specified in either of two ways: <ul style="list-style-type: none"> • 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. |
| <i>/prefix-length</i> | 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 Modes | |
|---------------|---------|
| | XR EXEC |

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | 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.

Use the **show route longer-prefixes** command to troubleshoot forwarding problems whose cause may be a long prefix.

| Task ID | Task ID | Operations |
|---------|---------|------------|
| | rib | read |

Examples

The following is sample output from the **show route longer-prefixes** command:

```
RP/0/RP0/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//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//CPU0/0
```

This table describes the significant fields shown in the display.

Table 131: 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 <i>ndnh</i>) since the route was installed in the RIB. |
| MgmtEth0//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 1125 | 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 XR EXEC mode.

```
show route [{ipv4 | ipv6}] [{unicast | {topology topo-name} | safi-all}] next-hop
[ip-address][{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. |
| safi-all | (Optional) Specifies unicast address prefixes. |
| <i>ip-address</i> | (Optional) IP address about which next-hop information is to be displayed. |
| standby | (Optional) Displays standby information. |

| Command Modes | XR EXEC |
|---------------|---------|
|---------------|---------|

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | 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.

Use the **show route next-hop** command to find all routes going through a next-hop address or interface.

| 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/RP0/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/RP0/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 132: 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 <i>ndzh</i>) 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 1101 | Displays the current contents of the routing table. |

show route quarantined

To display mutually recursive (looping) routes, use the **show route quarantined** command in XR EXEC mode.

```
show route [{ipv4 | ipv6}] [{unicast | {} | safi-all}] quarantined [{ip-address/prefix-length}ip-address mask}] [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. |
| safi-all | (Optional) Specifies unicast address prefixes. |
| <i>ip-address</i> | (Optional) IP address about which looping routes information is to be displayed. |
| <i>/ prefix-length</i> | (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. |
| <i>ip-address mask</i> | (Optional) Network mask applied to the <i>ip-address</i> argument. |
| standby | (Optional) Displays standby information. |

Command Modes XR EXEC

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | 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.

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.

| Task ID | Task ID | Operations |
|---------|---------|------------|
| | rib | read |

Examples

The following is sample output from the **show route quarantined** command:

```
RP/0/RP0/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 133: 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 <i>ndnh</i>) since the route was installed in the RIB. |
| (quarantined) | Shows that the route is quarantined. |

Related Commands

| Command | Description |
|------------------------------------------|-----------------------------------------------------|
| show route, on page 1101 | 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 XR EXEC mode.

```
show route [{ipv4 | ipv6}] [{unicast | {topology topo-name} | safi-all}] resolving-next-hop ip-address
[ 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. |
| safi-all | (Optional) Specifies unicast address prefixes. |
| <i>ip-address</i> | IP address about which resolved next-hop information is to be displayed. |
| standby | (Optional) Displays standby information. |

Command Modes

XR EXEC

Command History

| Release | Modification |
|---------------|------------------------------|
| 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.

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.

Task ID

| Task ID | Operations |
|---------|------------|
| rib | read |

Examples

The following is sample output from the **show route resolving-next-hop** command:

```
RP/0/RP0/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 134: 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 1101 | 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 XR EXEC mode.

```
show route [{afi-all | ipv4 | ipv6}] [{unicast | {topology topo-name} | safi-all}] static [ 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. |
| unicast | (Optional) Specifies unicast address prefixes. |
| safi-all | (Optional) Specifies unicast address prefixes. |
| standby | (Optional) Displays standby information. |

| Command Modes | |
|---------------|---------|
| | XR EXEC |

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | 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.

Use the **show route static** command to display information about static routes in the routing table.

| Task ID | Task ID | Operations |
|---------|---------|------------|
| | rib | read |

Examples

The following is sample output from the **show route static** command:

```
RP/0/RP0/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 135: 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 1101 | 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 XR EXEC mode.

```
show route [{afi-all | ipv4 | ipv6}] [{unicast | {topology topo-name} | safi-all}] summary [detail]
[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. |
| unicast | (Optional) Specifies unicast address prefixes. |
| safi-all | (Optional) Specifies unicast 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 Modes

XR EXEC

Command History

| Release | Modification |
|---------------|------------------------------|
| 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.

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.

Task ID

| Task ID | Operations |
|---------|------------|
| rib | read |

Examples

The following is sample output from the **show route summary** command:

```
RP/0/RP0/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 136: 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/RP0/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 137: 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. |

| Field | Description |
|--------------|-----------------------------------------------------------------------------|
| 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 1101 | Displays the current contents of the routing table. |

show route summary



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* module in the *Routing Configuration Guide for Cisco NCS 6000 Series Routers*.

- [authentication keychain mode \(RIP\), on page 1130](#)
- [auto-summary \(RIP\), on page 1132](#)
- [broadcast-for-v2, on page 1133](#)
- [default-information originate \(RIP\), on page 1134](#)
- [default-metric \(RIP\), on page 1135](#)
- [distance \(RIP\), on page 1136](#)
- [interface \(RIP\), on page 1138](#)
- [maximum-paths \(RIP\), on page 1140](#)
- [metric-zero-accept, on page 1141](#)
- [neighbor \(RIP\), on page 1142](#)
- [nsf \(RIP\), on page 1143](#)
- [output-delay, on page 1144](#)
- [passive-interface \(RIP\), on page 1145](#)
- [poison-reverse, on page 1146](#)
- [receive version, on page 1148](#)
- [redistribute \(RIP\), on page 1149](#)
- [router rip, on page 1152](#)
- [route-policy \(RIP\), on page 1154](#)
- [send version, on page 1155](#)
- [show protocols \(RIP\), on page 1156](#)
- [show rip, on page 1157](#)
- [show rip database, on page 1159](#)
- [show rip interface, on page 1161](#)
- [show rip statistics, on page 1164](#)
- [site-of-origin \(RIP\), on page 1166](#)
- [split-horizon disable \(RIP\), on page 1168](#)
- [timers basic, on page 1170](#)
- [validate-update-source disable, on page 1172](#)

authentication keychain mode (RIP)

To enable an authentication keychain mechanism on RIP interfaces, use the **authentication keychain mode** command in 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 | |
|----------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>keychain-name</i> | 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 <i>Implementing Keychain Management</i> module of <i>System Security Configuration Guide for Cisco NCS 6000 Series Routers</i> |
| 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

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | 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.

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 NCS 6000 Series Routers*

| Task ID | Task ID | Operation |
|---------|---------|----------------|
| | rip | read, write |

| Related Commands | Command | Description |
|------------------|------------------------------------------|---------------------------------------------------------------------------------------------------------------------|
| | router rip, on page 1152 | Configures a routing process and enters router configuration mode for a Routing Information Protocol (RIP) process. |

| Command | Description |
|------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| key chain (key chain) | Creates or modifies a keychain. Refer <i>System Security Command Reference for Cisco NCS 6000 Series Routers</i> for complete command reference information. |
| key (key chain) | Creates or modifies a keychain key. Refer <i>System Security Command Reference for Cisco NCS 6000 Series Routers</i> for complete command reference information. |
| key-string (keychain) | Specifies text string for the key. Refer <i>System Security Command Reference for Cisco NCS 6000 Series Routers</i> 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

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | 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.

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/RP0/CPU0:router(config)# router rip
RP/0/RP0/CPU0:router(config-rip)# auto-summary
```


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
 Interface configuration

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | 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 | 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/RP0/CPU0:router(config)# router rip
RP/0/RP0/CPU0:router(config-rip)# broadcast-for-v2
```

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 <i>name</i> Route policy name that indicates criteria for the default route. |
|---------------------------|--------------------------------------------------------------------------------------------------|

| | |
|------------------------|--------------------------------------|
| Command Default | This command is disabled by default. |
|------------------------|--------------------------------------|

| | |
|----------------------|----------------------|
| Command Modes | Router configuration |
|----------------------|----------------------|

| Command History | Release | Modification |
|------------------------|----------------|------------------------------|
| | 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 | 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: |
|-----------------|---------------------------------------------------------------------------------------------------------------------------------------------------|

```
RP/0/RP0/CPU0:router(config)# router rip
RP/0/RP0/CPU0:router(config-rip)# default-information originate route-policy policy1
```

| Related Commands | Command | Description |
|-------------------------|--------------------------------------------------|------------------------------------------------------------------------------------|
| | route-policy (RIP), on page 1154 | 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 | <i>number-value</i> Default metric value. Range is 1 to 15. |
|---------------------------|-------------------------------------------------------------|

| | |
|------------------------|------------------------------|
| Command Default | Default metrics are not set. |
|------------------------|------------------------------|

| | |
|----------------------|----------------------|
| Command Modes | Router configuration |
|----------------------|----------------------|

| | | |
|------------------------|----------------|------------------------------|
| Command History | Release | Modification |
| | 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.

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.

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 |

| | | |
|-------------------------|--------------------------------------------------|--------------------------------------------------------|
| Related Commands | Command | Description |
| | redistribute (RIP), on page 1149 | 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

| | |
|-----------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>admin-distance</i> | Administrative distance to be assigned to RIP routes. Range is 0 to 255. |
| <i>prefix</i> | (Optional) Network IP address about which routing information should be displayed. |
| <i>prefix-length</i> | (Optional) The <i>prefix-length</i> 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 must precede the decimal value. Range is 0 to 32 for IPv4 addresses. |
| <i>mask</i> | (Optional) Network mask specified in either of two ways: <ul style="list-style-type: none"> • 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

Command History

| Release | Modification |
|---------------|------------------------------|
| 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.

Use the **distance** command to change the preference of RIP routes over other protocol routes. When administrative distance and redistribution features are used 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 138: 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 prefix:

```
RP/0/RP0/CPU0:router(config)# router rip
RP/0/RP0/CPU0:router(config-rip)# distance 85 192.168.10.0/24
```

Related Commands

| Command | Description |
|--------------------------------------------------|--------------------------------------------------------|
| redistribute (RIP), on page 1149 | 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 | <i>type</i> | Interface type. For more information, use the question mark (?) online help function. |
| | <i>interface-path-id</i> | 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 |
|----------------------|----------------------|

| | | |
|------------------------|----------------|------------------------------|
| Command History | Release | Modification |
| | 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. |
|-------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

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 1133 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/RP0/CPU0:router(config)# router rip
RP/0/RP0/CPU0:router(config-rip)# interface GigabitEthernet 0/1/0/0
```

```

RP/0/RP0/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/RP0/CPU0:router(config-rip-if)# broadcast-for-v2

```

Related Commands

| Command | Description |
|------------------------------------------------|------------------------------------------------------------|
| broadcast-for-v2, on page 1133 | 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

Command History

Release

Release 5.0.0

Modification

This command was introduced.

Usage Guidelines

To use this command, you must be in a user group associated with a task 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/RP0/CPU0:router(config)# router rip
RP/0/RP0/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 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.

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:

```
RP/0/RP0/CPU0:router(config)# router rip
RP/0/RP0/CPU0:router(config-rip)# interface GigabitEthernet 0/1/0/0
RP/0/RP0/CPU0:router(config-rip-if)# metric-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 *ip-address*
no neighbor *ip-address*

| | |
|---------------------------|--------------------------------------------------------------------------------------------|
| Syntax Description | <i>ip-address</i> IP address of a peer router with which routing information is exchanged. |
|---------------------------|--------------------------------------------------------------------------------------------|

| | |
|------------------------|-------------------------------------|
| Command Default | No neighboring routers are defined. |
|------------------------|-------------------------------------|

| | |
|----------------------|----------------------|
| Command Modes | Router configuration |
|----------------------|----------------------|

| Command History | Release | Modification |
|------------------------|----------------|------------------------------|
| | Release 5.0.0 | This command was introduced. |

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

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 |

| Related Commands | Command | Description |
|-------------------------|-------------------------------------------------------|--------------------------------------------------------|
| | passive-interface (RIP), on page 1145 | 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 nsf** 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

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | 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.

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 |

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 | <i>delay</i> 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 | |
| Command History | Release | Modification |
| | 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. | |
| | 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 |

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 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 | Operations |
|---------|---------|----------------|
| | rip | read, write |

| Related Commands | Command | Description |
|------------------|----------------------------------------------|-------------------------------------------------------------------------------|
| | neighbor (RIP), on page 1142 | 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 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.

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:

```
RP/0/RP0/CPU0:router(config)# router rip
RP/0/RP0/CPU0:router(config-rip)# interface GigabitEthernet 0/1/0/0
RP/0/RP0/CPU0:router(config-rip-if)# poison-reverse
```

Related Commands

| Command | Description |
|-----------------------------------------------------------|---------------------------------------|
| split-horizon disable (RIP), on page 1168 | 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 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.

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/RP0/CPU0:router(config)# router rip
RP/0/RP0/CPU0:router(config-rip)# interface GigabitEthernet 0/1/0/0
RP/0/RP0/CPU0:router(config-rip-if)# receive version 1 2
```

| Related Commands | Command | Description |
|------------------|--------------------------------------------|----------------------------------------------------------------|
| | send version, on page 1155 | 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 | <ul style="list-style-type: none"> For the bgp keyword: <ul style="list-style-type: none"> 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. 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 <i>name</i> | (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 <i>name</i>] | <p>(Optional) Specifies the criteria by which OSPF routes are redistributed into other routing domains. It can be one or more of the following:</p> <ul style="list-style-type: none"> 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. <p>For the external and nssa-external options, if a type is not specified, then both Type 1 and Type 2 are assumed.</p> <p>If no match is specified, the default is no filtering.</p> |
| static | Redistributes IP static routes. |

Command Default Route redistribution is disabled.

Command Modes Router configuration

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | 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.



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.

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* module of the *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/RP0/CPU0:router(config)# router rip
RP/0/RP0/CPU0:router(config-rip)# redistribute bgp 100
```

| Related Commands | Command | Description |
|------------------|----------------------------------------------------|------------------------------------------------------------------------------------|
| | default-metric (RIP), on page 1135 | 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 XR Config 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 XR Config

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | 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 | Operations |
|---------|---------|----------------|
| | rip | read, write |

Examples

The following example shows how to configure a router process for RIP:

```
RP/0/RP0/CPU0:router(config)# router rip
RP/0/RP0/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/RP0/CPU0:router(config)# router rip
RP/0/RP0/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
RP/0/RP0/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
Interface configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| 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.

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 |
|---------|----------------|
| rip | read, write |

Examples

The following example shows how to filter routing updates received on an interface:

```
RP/0/RP0/CPU0:router(config)# router rip
RP/0/RP0/CPU0:router(config-rip)# interface GigabitEthernet 0/1/0/0
RP/0/RP0/CPU0: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 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.

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/RP0/CPU0:router(config)# router rip
RP/0/RP0/CPU0:router(config-rip)# interface GigabitEthernet 0/1/0/0
RP/0/RP0/CPU0:router(config-rip-if)# send version 2
```

| Related Commands | Command | Description |
|------------------|-----------------------------------------------|------------------------------------------------------------------|
| | receive version, on page 1148 | 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 XR EXEC mode.

```
show protocols [{ipv4 | afi-all}] [{allprotocol}] [{default-context}] [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. |
| <i>protocol</i> | (Optional) Specifies a routing protocol. <ul style="list-style-type: none"> • 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. |
| private | (Optional) Displays private EIGRP data. This keyword is available when the eigrp protocol is specified. |

| Command Modes | |
|---------------|---------|
| | XR EXEC |

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | 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.

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 |

show rip

To display configuration and status of Routing Information Protocol (RIP), use the **show rip** command in XR EXEC mode.

show rip

Command Default No default behavior or values

Command Modes XR EXEC

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | 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 | Operations |
|---------|---------|------------|
| | rip | read |

Examples

The following example shows sample output from the **show rip** command:

```
RP/0/RP0/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 139: show rip Field Descriptions

| Field | Description |
|---------|-----------------------|
| Active? | Active state setting. |

| Field | Description |
|---------------------------|----------------------------------------------------------------------------------------------|
| 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 XR EXEC mode.

```
show rip database [{prefix prefix-length | prefix mask}]
```

| Syntax Description | |
|----------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>prefix</i> | (Optional) Network IP address about which routing information should be displayed. |
| <i>prefix-length</i> | (Optional) The <i>prefix-length</i> 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. |
| <i>prefix-mask</i> | (Optional) Network mask specified in either of two ways: <ul style="list-style-type: none"> • 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 | XR EXEC |
|---------------|---------|
|---------------|---------|

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | 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.

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 database** command:

```
RP/0/RP0/CPU0:router# show rip database

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

```

This table describes the significant fields shown in the display.

Table 140: show rip 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. |
| 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. |

show rip interface

To display interface entry information from the Routing Information Protocol (RIP) topology table, use the **show rip interface** command in XR EXEC mode.

```
show rip interface [type interface-path-id]
```

| Syntax Description | interface | (Optional) Specifies the interface from which to clear topology entries. |
|--------------------|--------------------------|------------------------------------------------------------------------------------------------------------|
| | <i>type</i> | Interface type. For more information, use the question mark (?) online help function. |
| | <i>interface-path-id</i> | 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 | XR EXEC |
|----------------------|---------|

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | 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 | Operations |
|---------|---------|------------|
| | rip | read |

Examples

This example is sample output from the **show rip interface** command:

```
RP/0/RP0/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

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

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 141: 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. |
| 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. |

| Field | Description |
|-----------------------|------------------------------------------------------------------|
| 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 XR EXEC mode.

show rip statistics

Command Default No default behavior or values

Command Modes XR EXEC

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | 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 | Operations |
|---------|---------|------------|
| | rip | read |

Examples

The following example is sample output from the **show rip statistics** command:

```
RP/0/RP0/CPU0:router# show rip statistics

RIP statistics:
  Total messages sent:          5597
  Message send failures:       0
  Regular updates sent:        5566
  Queries responded 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
```

This table describes the significant fields shown in the display.

Table 142: show rip statistics Field Descriptions

| Field | Description |
|---------------------|-----------------------------|
| Total messages sent | Number of RIP packets sent. |

| Field | Description |
|----------------------------|-----------------------------------------------------------------------|
| Message send failures | Number of times that the packet send operation failed. |
| Queries responded 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 not 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) in 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 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.

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/RP0/CPU0:router(config)# router rip  
RP/0/RP0/CPU0:router(config-rip) interface GigabitEthernet 0/1/0/0  
RP/0/RP0/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 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.

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:

```
RP/0/RP0/CPU0:router(config)# router rip
RP/0/RP0/CPU0:router(config-rip)# interface GigabitEthernet 0/1/0/0
RP/0/RP0/CPU0:router(config-rip-if)# split-horizon disable
```

Related Commands

| Command | Description |
|----------------------------------------------|----------------------------------------------------------|
| poison-reverse, on page 1146 | 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

| | |
|-----------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>update</i> | Rate, in seconds, at which updates are sent. This is the fundamental timing parameter of the routing protocol. Range is 5 to 50000. |
| <i>invalid</i> | 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. |
| <i>holddown</i> | 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. |
| <i>flush</i> | 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 <i>invalid</i> 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 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.

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/RP0/CPU0:router(config)# router rip  
RP/0/RP0/CPU0:router(config-rip) timers basic 5 15 15 30
```

Related Commands

| Command | Description |
|----------------------------------------|-------------------------------------------|
| show rip, on page 1157 | 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 reenble 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 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.

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/RP0/CPU0:router(config)# router rip
RP/0/RP0/CPU0:router(config-rip) validate-update-source disable
```




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* module in the *Routing Configuration Guide for Cisco NCS 6000 Series Routers*.

- [abort \(RPL\), on page 1177](#)
- [add, on page 1179](#)
- [apply, on page 1181](#)
- [as-path in, on page 1183](#)
- [as-path is-local, on page 1185](#)
- [as-path length, on page 1187](#)
- [as-path neighbor-is, on page 1189](#)
- [as-path originates-from, on page 1192](#)
- [as-path passes-through, on page 1194](#)
- [as-path-set, on page 1196](#)
- [as-path unique-length, on page 1198](#)
- [community is-empty, on page 1200](#)
- [community matches-any, on page 1201](#)
- [community matches-every, on page 1203](#)
- [community matches-within, on page 1205](#)
- [community-set, on page 1207](#)
- [delete community, on page 1210](#)
- [delete large-community, on page 1212](#)
- [delete extcommunity rt, on page 1214](#)
- [destination in, on page 1216](#)
- [done, on page 1218](#)
- [drop, on page 1220](#)
- [edit, on page 1222](#)
- [end-global, on page 1225](#)
- [end-policy, on page 1226](#)
- [end-set, on page 1227](#)
- [extcommunity rt is-empty, on page 1229](#)
- [extcommunity rt matches-any, on page 1230](#)
- [extcommunity rt matches-every, on page 1232](#)

- `extcommunity rt matches-within`, on page 1234
- `extcommunity-set cost`, on page 1236
- `extcommunity-set rt`, on page 1238
- `extcommunity-set soo`, on page 1240
- `extcommunity soo is-empty`, on page 1242
- `extcommunity soo matches-any`, on page 1243
- `extcommunity soo matches-every`, on page 1245
- `if`, on page 1247
- `local-preference`, on page 1253
- `large-community is-empty`, on page 1254
- `large-community matches-any`, on page 1256
- `large-community matches-every`, on page 1258
- `large-community matches-within`, on page 1260
- `large-community-set`, on page 1262
- `med`, on page 1263
- `next-hop in`, on page 1264
- `orf prefix in`, on page 1266
- `origin is`, on page 1268
- `ospf-area-set`, on page 1270
- `pass`, on page 1272
- `path-type is`, on page 1274
- `policy-global`, on page 1275
- `prefix-set`, on page 1277
- `prepend as-path`, on page 1279
- `protocol`, on page 1281
- `rd in`, on page 1283
- `replace as-path`, on page 1284
- `rib-has-route`, on page 1286
- `route-has-label`, on page 1288
- `route-policy (RPL)`, on page 1289
- `route-type is`, on page 1291
- `rpl editor`, on page 1293
- `rpl maximum`, on page 1294
- `set aigp-metric`, on page 1296
- `set community`, on page 1297
- `set dampening`, on page 1299
- `set eigrp-metric`, on page 1301
- `set extcommunity cost`, on page 1303
- `set extcommunity rt`, on page 1305
- `set ip-precedence`, on page 1307
- `set isis-metric`, on page 1309
- `set label`, on page 1310
- `set large-community`, on page 1311
- `set level`, on page 1313
- `set local-preference`, on page 1314
- `set med`, on page 1315

- [set metric-type \(IS-IS\), on page 1317](#)
- [set metric-type \(OSPF\), on page 1318](#)
- [set next-hop, on page 1319](#)
- [set origin, on page 1321](#)
- [set ospf-metric, on page 1322](#)
- [set path-selection, on page 1323](#)
- [set qos-group \(RPL\), on page 1325](#)
- [set rib-metric, on page 1326](#)
- [set rip-metric, on page 1327](#)
- [set rip-tag, on page 1328](#)
- [set rtset route-limit , on page 1329](#)
- [set spf-priority, on page 1330](#)
- [set tag, on page 1331](#)
- [set traffic-index, on page 1332](#)
- [set weight, on page 1334](#)
- [show rpl, on page 1336](#)
- [show rpl active as-path-set, on page 1338](#)
- [show rpl active community-set, on page 1341](#)
- [show rpl active extcommunity-set, on page 1344](#)
- [show rpl active prefix-set, on page 1347](#)
- [show rpl active rd-set, on page 1350](#)
- [show rpl active route-policy, on page 1352](#)
- [show rpl as-path-set, on page 1354](#)
- [show rpl as-path-set attachpoints, on page 1356](#)
- [show rpl as-path-set references, on page 1359](#)
- [show rpl community-set, on page 1361](#)
- [show rpl community-set attachpoints, on page 1363](#)
- [show rpl community-set references, on page 1365](#)
- [show rpl extcommunity-set, on page 1368](#)
- [show rpl inactive as-path-set, on page 1371](#)
- [show rpl inactive community-set, on page 1374](#)
- [show rpl inactive extcommunity-set, on page 1377](#)
- [show rpl inactive prefix-set, on page 1380](#)
- [show rpl inactive rd-set, on page 1383](#)
- [show rpl inactive route-policy, on page 1385](#)
- [show rpl maximum, on page 1388](#)
- [show rpl policy-global references, on page 1390](#)
- [show rpl prefix-set, on page 1392](#)
- [show rpl prefix-set attachpoints, on page 1394](#)
- [show rpl prefix-set references, on page 1397](#)
- [show rpl rd-set, on page 1399](#)
- [show rpl rd-set references, on page 1401](#)
- [show rpl route-policy, on page 1403](#)
- [show rpl route-policy attachpoints, on page 1406](#)
- [show rpl route-policy inline, on page 1408](#)
- [show rpl route-policy references, on page 1410](#)

- [show rpl route-policy uses](#), on page 1413
- [show rpl unused as-path-set](#), on page 1416
- [show rpl unused community-set](#), on page 1419
- [show rpl unused extcommunity-set](#), on page 1422
- [show rpl unused prefix-set](#), on page 1424
- [show rpl unused rd-set](#), on page 1427
- [show rpl unused route-policy](#), on page 1429
- [source in](#), on page 1432
- [source rt](#), on page 1434
- [suppress-route](#), on page 1435
- [tag](#), on page 1436
- [tag in](#), on page 1437
- [tag-set](#), on page 1439
- [unsuppress-route](#), on page 1440

abort (RPL)

To discard a route policy or set definition and return to XR Config 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 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 | Operations |
|--------------|----------------|
| route-policy | read, write |

Examples

The following example shows how to discard a route policy definition that was started and return to XR Config mode:

```
RP/0/RP0/CPU0:router(config)# route-policy policy_1
RP/0/RP0/CPU0:router(config-rpl)# if as-path is-local then
RP/0/RP0/CPU0:router(config-rpl-if)# abort
RP/0/RP0/CPU0:router(config)#
```

The following example shows how to discard a prefix set definition that was started and return to XR Config mode:

```
RP/0/RP0/CPU0:router(config)# prefix-set legal-ipv4-prefix-examples  
RP/0/RP0/CPU0:router(config-pfx)# 10.0.1.1,  
RP/0/RP0/CPU0:router(config-pfx)# 10.0.2.0/24,  
RP/0/RP0/CPU0:router(config-pfx)# abort  
RP/0/RP0/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. | |
| <i>bandwidth</i> | Bandwidth in kilobits per second. Range is from 0 to 4294967295. | |
| <i>delay</i> | Delay in 10-microsecond units. Range is from 0 to 4294967295. | |
| <i>reliability</i> | Reliability metric. 255 is 100-percent reliable. Range is from 0 to 255. | |
| <i>loading</i> | Effective bandwidth (loading). 255 is 100-percent loaded. Range is from 0 to 255. | |
| <i>max-transmission</i> | Maximum transmission of the path. Range is from 0 to 65535. | |
| rip-metric | Specifies an RIP metric attribute. | |
| <i>number</i> | Value assigned to a four-bit unsigned integer. Range is from 0 to 16. | |
| <i>parameter</i> | 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 5.0.0 | This command was introduced. |
| Usage Guidelines | <p>To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.</p> <p>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).</p> | |
| Task ID | Task ID | Operations |
| | route-policy | read, write |

Examples

The following example shows how to offset the RIP metric value:

```
RP/0/RP0/CPU0:router(config)# route-policy policy_1
```

```
RP/0/RP0/CPU0:router(config-rpl)# add rip-metric 4  
RP/0/RP0/CPU0:router(config-rpl)# end-policy
```

The following example shows how to set the EIGRP metric value:

```
RP/0/RP0/CPU0:router(config)# route-policy policy_1  
RP/0/RP0/CPU0:router(config-rpl)# add eigrp-metric 50000 24000 230 14000  
RP/0/RP0/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 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.

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.

Wildcards can be used in apply policy names. This supports the nested wildcard apply scenario. A wildcard is specified by inserting an asterisk (*) in place of one of the portions of the apply policy name; the wildcard indicates that any value for that portion of the apply policy name matches. The nested wildcard apply policy allows wildcard (*) based apply nesting. The wildcard operation permits declaration of a generic apply statement that calls all policies that contain a specific defined set of alphanumeric characters, defined on the router.

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/RP0/CPU0:router(config)# route-policy SetLocalPref ($lp0, $lp1, $lp2, $lp3, $lp4)
RP/0/RP0/CPU0:router(config-rpl)# if community matches-any ($lp0:$lp1) then
RP/0/RP0/CPU0:router(config-rpl-elseif)# set local-preference $lp1
RP/0/RP0/CPU0:router(config-rpl-elseif)# elseif community matches-any ($lp0:$lp2) then
RP/0/RP0/CPU0:router(config-rpl-elseif)# set local-preference $lp2
RP/0/RP0/CPU0:router(config-rpl-elseif)# elseif community matches-any ($lp0:$lp3) then
RP/0/RP0/CPU0:router(config-rpl-elseif)# set local-preference $lp3
RP/0/RP0/CPU0:router(config-rpl-elseif)# elseif community matches-any ($lp0:$lp4) then
RP/0/RP0/CPU0:router(config-rpl-elseif)# set local-preference $lp4
RP/0/RP0/CPU0:router(config-rpl-elseif)# endif
RP/0/RP0/CPU0:router(config-rpl)# end-policy

RP/0/RP0/CPU0:router(config)# route-policy CustomerIn($cust)
RP/0/RP0/CPU0:router(config-rpl)# apply SetLocalPref ($cust, 20, 30, 40, 50)
RP/0/RP0/CPU0:router(config-rpl)# end-policy

RP/0/RP0/CPU0:router(config)# route-policy Cust_217
RP/0/RP0/CPU0:router(config-rpl)# apply CustomerIn(217)
RP/0/RP0/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-name*inline-*as-path-set-parameter*}

| Syntax Description | |
|---------------------------|-----------------------------------------------------------------------------|
| <i>as-path-set-name</i> | Name of an AS path set. |
| <i>inline-as-path-set</i> | Inline AS path set. The inline AS path set must be enclosed in parentheses. |
| <i>parameter</i> | 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 3.9.0 | No modification. |
| | 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.

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/RP0/CPU0:router(config)# as-path-set my-as-set
RP/0/RP0/CPU0:router(config-as)# ios-regex '_12$',
```

```
RP/0/RP0/CPU0:router(config-as)# ios-regex '_13$'  
RP/0/RP0/CPU0:router(config-as)# end-set
```

and the following policy excerpt using an *as-path-set-name* argument:

```
RP/0/RP0/CPU0:router(config-rpl)# if as-path in my-as-set then  
RP/0/RP0/CPU0:router(config-rpl-if)# set local-preference 100  
RP/0/RP0/CPU0:router(config-rpl-if)# endif  
RP/0/RP0/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/RP0/CPU0:router(config-rpl)# if as-path in (ios-regex '_12$',ios-regex '_13$') then  
RP/0/RP0/CPU0:router(config-rpl-if)# set local-preference 100  
RP/0/RP0/CPU0:router(config-rpl-if)# endif  
RP/0/RP0/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 3.9.0 | No modification. |
| | 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.

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/RP0/CPU0:router(config-rpl)# if as-path is-local then
RP/0/RP0/CPU0:router(config-rpl-if)# set local-preference 100
```

```
RP/0/RP0/CPU0:router(config-rpl-if)# endif  
RP/0/RP0/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. |
|--------------------|-------------------|------------------------------------------------------------------------|
| | <i>number</i> | Value assigned to an 11-bit unsigned integer. Range is from 0 to 2047. |
| | <i>parameter</i> | 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 3.9.0 | No modification. |
| | 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.

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/RP0/CPU0:router(config-rpl)# if as-path length eq 10 then
RP/0/RP0/CPU0:router(config-rpl-if)# set local-preference 100
RP/0/RP0/CPU0:router(config-rpl-if)# endif
RP/0/RP0/CPU0:router(config-rpl)#
```

Related Commands

| Command | Description |
|-------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| as-path in, on page 1183 | Matches the AS path of a route to an AS path set. |
| as-path originates-from, on page 1192 | Compares an AS path against the AS sequence beginning with the AS number that originated a route |
| as-path passes-through, on page 1194 | 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 1198 | 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) in 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 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.

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/RP0/CPU0:router(config-rpl)# if as-path neighbor-is '10' then
RP/0/RP0/CPU0:router(config-rpl-if)# if as-path neighbor-is '$asnum' then
RP/0/RP0/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/RP0/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/RP0/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 1183 | Matches the AS path of a route to an AS path set. |
| as-path length, on page 1187 | Compares the number of ASN in the AS path of a route. |
| as-path originates-from, on page 1192 | Compares an AS path to the AS sequence beginning with the AS number that originated a route. |
| as-path passes-through, on page 1194 | 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 1198 | 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) in 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 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.

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/RP0/CPU0:router(config-rpl)# if as-path originates-from '10 11' then
RP/0/RP0/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/RP0/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/RP0/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 1183 | Matches the AS path of a route to an AS path set. |
| | as-path length, on page 1187 | Compares the number of ASN in the AS path of a route |
| | as-path passes-through, on page 1194 | 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 1198 | 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) in 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 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.

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/RP0/CPU0:router(config-rpl)# if as-path passes-through '10' then
RP/0/RP0/CPU0:router(config-rpl-if)# if as-path passes-through '$asnum' then
RP/0/RP0/CPU0:router(config-rpl-if)# if as-path passes-through '10 11' then
RP/0/RP0/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/RP0/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/RP0/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 1183 | Matches the AS path of a route to an AS path set. |
| | as-path length, on page 1187 | Compares the number of ASN in the AS path of a route |
| | as-path originates-from, on page 1192 | Compares an AS path to the AS sequence beginning with the AS number that originated a route. |
| | as-path unique-length, on page 1198 | Performs specific checks based on the length of the AS path. |

as-path-set

To create a named AS path set, use the **as-path-set** command in XR Config mode. To remove the named AS path set, use the **no** form of this command.

as-path-set *name*
no as-path-set *name*

| | |
|---------------------------|--------------------------------------|
| Syntax Description | <i>name</i> Name of the AS path set. |
|---------------------------|--------------------------------------|

| | |
|------------------------|-------------------------------|
| Command Default | No default behavior or values |
|------------------------|-------------------------------|

| | |
|----------------------|-----------|
| Command Modes | XR Config |
|----------------------|-----------|

| Command History | Release | Modification |
|------------------------|----------------|------------------------------|
| | Release 3.9.0 | No modification. |
| | 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.

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. |
| 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. |

| Options | Description |
|----------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 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/RP0/CPU0:router(config)# as-path-set aset1
RP/0/RP0/CPU0:router(config-as)# ios-regex '_42$',
RP/0/RP0/CPU0:router(config-as)# ios-regex '_127$'
RP/0/RP0/CPU0:router(config-as)# end-set
```

The following is a sample of the as-path options used as an inline set.

```
RP/0/RP0/CPU0:router(config-rpl)# if as-path in (ios-regex '_42$', ios-regex$ '_127$')
RP/0/RP0/CPU0:router(config-rpl-if)# pass
RP/0/RP0/CPU0:router(config-rpl-if)# endif
RP/0/RP0/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. |
| <i>number</i> | Value assigned to an 11-bit unsigned integer. Range is from 0 to 2047. |
| <i>parameter</i> | 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 3.9.0 | No modification. |
| | 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.

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 how 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/RP0/CPU0:router(config-rpl)# if as-path unique-length eq 10 then
```

```
RP/0/RP0/CPU0:router(config-rpl-if)# if as-path unique-length ge 10 then
RP/0/RP0/CPU0:router(config-rpl-if)# if as-path unique-length le 10 then

RP/0/RP0/CPU0:router(config-rpl)# if as-path unique-length eq $integerparam then
RP/0/RP0/CPU0:router(config-rpl-if)# if as-path unique-length ge $geparam then
RP/0/RP0/CPU0:router(config-rpl-if)# if as-path unique-length le $leparam then

RP/0/RP0/CPU0:router(config-rpl)# set local-preference 100
RP/0/RP0/CPU0:router(config-rpl)# endif
```

Related Commands

| Command | Description |
|----------------------------------------------|-----------------------------------------------------------------|
| as-path length, on page 1187 | 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 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.

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/RP0/CPU0:router(config-rpl)# if community is-empty then
RP/0/RP0/CPU0:router(config-rpl-if)# set local-preference 100
RP/0/RP0/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-name*inline-*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

Release 5.0.0

Modification

This command was introduced.

Usage Guidelines

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

community matches-any

```
[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/RP0/CPU0:router(config)# community-set my-community-set
RP/0/RP0/CPU0:router(config-comm)# 10:20,
RP/0/RP0/CPU0:router(config-comm)# 10:30,
RP/0/RP0/CPU0:router(config-comm)# 10:40
RP/0/RP0/CPU0:router(config-comm)# end-set

RP/0/RP0/CPU0:router(config)# route-policy community-matches-any-example
RP/0/RP0/CPU0:router(config-rpl)# if community matches-any my-community-set then
RP/0/RP0/CPU0:router(config-rpl-if)# set local-preference 100
RP/0/RP0/CPU0:router(config-rpl-if)# elseif community matches-any ([10..25]:35) then
RP/0/RP0/CPU0:router(config-rpl-elseif)# set local-preference 200
RP/0/RP0/CPU0:router(config-rpl-elseif)# elseif community matches-any (30:[100..500])
then
RP/0/RP0/CPU0:router(config-rpl-elseif)# set local-preference 300
RP/0/RP0/CPU0:router(config-rpl-elseif)# endif
RP/0/RP0/CPU0:router(config-rpl)# end-policy
```

Related Commands

| Command | Description |
|-------------------------------------------------------|-------------------------------------------|
| community matches-every, on page 1203 | 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

Release 5.0.0

Modification

This command was introduced.

Usage Guidelines

To use this command, you must be in a user group associated with a task group that 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/RP0/CPU0:router(config)# community-set my-community-set
RP/0/RP0/CPU0:router(config-comm)# 10:20,
RP/0/RP0/CPU0:router(config-comm)# 10:30,
RP/0/RP0/CPU0:router(config-comm)# 10:40
RP/0/RP0/CPU0:router(config-comm)# end-set

RP/0/RP0/CPU0:router(config)# route-policy community-matches-every-example
RP/0/RP0/CPU0:router(config-rpl)# if community matches-every my-community-set then
RP/0/RP0/CPU0:router(config-rpl-if)# set local-preference 100
RP/0/RP0/CPU0:router(config-rpl-elseif)# elseif community matches-every (ios-regex
' _10:[0-9]0_') then
RP/0/RP0/CPU0:router(config-rpl-elseif)# set local-preference 200
RP/0/RP0/CPU0:router(config-rpl-elseif)# elseif community matches-every
(ios-regex' _20:[0-9]0_') then
RP/0/RP0/CPU0:router(config-rpl-elseif)# set local-preference 300
RP/0/RP0/CPU0:router(config-rpl-elseif)# endif
RP/0/RP0/CPU0:router(config-rpl)# end-policy
```

Related Commands

| Command | Description |
|-----------------------------------------------------|-----------------------------------------|
| community matches-any, on page 1201 | 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/RP0/CPU0:router#config
RP/0/RP0/CPU0:router(config)#route-policy bob
RP/0/RP0/CPU0:router(config-rpl)#if community matches-within (*:3, 5:*) then
RP/0/RP0/CPU0:router(config-rpl)#set local-preference 94
```

```
RP/0/RP0/CPU0:router (config-rpl) #endif  
RP/0/RP0/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 XR Config mode. To remove the community set, use the **no** form of this command.

```
community-set name
no community-set name
```

| | |
|---------------------------|----------------------------------------|
| Syntax Description | <i>name</i> Name of the community set. |
|---------------------------|----------------------------------------|

| | |
|------------------------|-------------------------------|
| Command Default | No default behavior or values |
|------------------------|-------------------------------|

| | |
|----------------------|-----------|
| Command Modes | XR Config |
|----------------------|-----------|

| Command History | Release | Modification |
|------------------------|----------------|------------------------------|
| | 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.

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: **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 |
|-----------------|------------------------------------------|
| <i>#-remark</i> | Remark beginning with '#' |
| * | Wildcard (any community or part thereof) |
| <i>0-65535</i> | 16-bit half-community number |

| Format | Description |
|---------------------|--------------------------------------------------------------|
| [| Left bracket to begin range |
| 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 cset1 is created:

```
RP/0/RP0/CPU0:router(config)# community-set cset1
RP/0/RP0/CPU0:router(config-comm)# 12:34,
RP/0/RP0/CPU0:router(config-comm)# 12:56,
RP/0/RP0/CPU0:router(config-comm)# 12:78,
RP/0/RP0/CPU0:router(config-comm)# internet
RP/0/RP0/CPU0:router(config-comm)# end-set
```

In the following example, a community set named cset2 is created:

```
RP/0/RP0/CPU0:router(config)# community-set cset2
RP/0/RP0/CPU0:router(config-comm)# 123:456,
RP/0/RP0/CPU0:router(config-comm)# no-advertise,
RP/0/RP0/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/RP0/CPU0:router(config)# community-set cset3  
RP/0/RP0/CPU0:router(config-comm)# 123:*  
RP/0/RP0/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. |
| <i>community-set-name</i> | Name of a community set. |
| <i>inline-community-set</i> | Inline community set. The inline community set must be enclosed in parentheses. |
| <i>parameter</i> | 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 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.

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/RP0/CPU0:router(config-rpl)# delete community in my_community_set
RP/0/RP0/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/RP0/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/RP0/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/RP0/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. |
| | <i>large-community-set-name</i> | Name of a large community set. |
| | <i>inline-large-community-set</i> | Inline large community set. The inline community set must be enclosed in parentheses. |
| | <i>parameter</i> | 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.
- peers — This expression is replaced by the AS number of the neighbor 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 peers.
- private-as — This expression specifies any number in the private ASN range: [64512..65534] and [4200000000..4294967294].



Note The `peer-as` and `not-peer-as` 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/RP0/RP0/CPU0:router#config
RP/0/RP0/CPU0:router(config)#route-policy lrg_comm_rp2S
RP/0/RP0/CPU0:router(config-rpl)#delete large-community in (ios-regex '^100000:')
RP/0/RP0/CPU0:router(config-rpl)#delete large-community not in (peer-as:*, 41289:*)
RP/0/RP0/CPU0:router(config-rpl)#delete large-community in catbert
RP/0/RP0/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. |
| <i>extcommunity-set-name</i> | Name of an extended community set. |
| <i>inline-extcommunity-set</i> | Inline extended community set. The inline extended community set must be enclosed in parentheses. |
| <i>parameter</i> | 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

Usage Guidelines

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 5.0.0 | 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 **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/RP0/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/RP0/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/RP0/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-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 name. The parameter name must be preceded with a "\$."

parameter

Command Default

No default behavior or values

Command Modes

Route-policy configuration

Command History

Release

Release 5.0.0

Modification

This command was introduced.

Usage Guidelines

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

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.



Note

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.

In Border Gateway Protocol (BGP), the destination of a route is also known as its network-layer reachability information (NLRI). It comprises a prefix value and a mask length.

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/RP0/CPU0:router(config)# prefix-set my-prefix-set
RP/0/RP0/CPU0:router(config-pfx)# 10.0.0.1/32,
RP/0/RP0/CPU0:router(config-pfx)# fe80::203:0:0:0/64,
RP/0/RP0/CPU0:router(config-pfx)# 10.0.0.2/24 le 32
RP/0/RP0/CPU0:router(config-pfx)# end-set

RP/0/RP0/CPU0:router(config)# route-policy use-destination-in
RP/0/RP0/CPU0:router(config-rpl)# if destination in my-prefix-set then
RP/0/RP0/CPU0:router(config-rpl-if)# set local-preference 100
RP/0/RP0/CPU0:router(config-rpl-if)# elseif destination in (10.0.0.1/32, 10.0.0.2/24 le
32) then
RP/0/RP0/CPU0:router(config-rpl-elseif)# set local-preference 200
RP/0/RP0/CPU0:router(config-rpl-elseif)# endif
RP/0/RP0/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/RP0/CPU0:router(config)# prefix-set ipv6-prefix-set
RP/0/RP0/CPU0:router(config-pfx)# 2001:0:0:1::/64,
RP/0/RP0/CPU0:router(config-pfx)# 2001:0:0:2::/64,
RP/0/RP0/CPU0:router(config-pfx)# 2001:0:0:3::/64,
RP/0/RP0/CPU0:router(config-pfx)# 2001:0:0:4::/64
RP/0/RP0/CPU0:router(config-pfx)# end-set

RP/0/RP0/CPU0:router(config)# route-policy ipv6-destination-in
RP/0/RP0/CPU0:router(config-rpl)# if destination in ipv6-prefix-set then
RP/0/RP0/CPU0:router(config-rpl-if)# set next-hop 2001:abcd:fedc::1
RP/0/RP0/CPU0:router(config-rpl-if)# elseif destination in (2001::1, 2002:1:2:3::/64)
then
RP/0/RP0/CPU0:router(config-rpl-elseif)# set next-hop
1111:2222:3333:4444:5555:6666:7777:8888
RP/0/RP0/CPU0:router(config-rpl-elseif)# endif
RP/0/RP0/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 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.

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 it encounters the *done* statement.

```
RP/0/RP0/CPU0:router(config)# route-policy done_st_example
RP/0/RP0/CPU0:router(config-rpl)# if destination in (29.0.0.0/8 le 32) then
```

```
RP/0/RP0/CPU0:router(config-rpl-if)# set community 102:12
RP/0/RP0/CPU0:router(config-rpl-if)# endif
RP/0/RP0/CPU0:router(config-rpl)# if destination in (39.0.0.0/8 le 32) then
RP/0/RP0/CPU0:router(config-rpl-if)# set community 102:39
RP/0/RP0/CPU0:router(config-rpl-if)# done
RP/0/RP0/CPU0:router(config-rpl-if)# endif
RP/0/RP0/CPU0:router(config-rpl)# if destination in (49.0.0.0/8 le 32) then
RP/0/RP0/CPU0:router(config-rpl-if)# set community 102:49
RP/0/RP0/CPU0:router(config-rpl-if)# endif
RP/0/RP0/CPU0:router(config-rpl)# if destination in (59.0.0.0/8 le 32) then
RP/0/RP0/CPU0:router(config-rpl-if)# set community 102:59
RP/0/RP0/CPU0:router(config-rpl-if)# endif
RP/0/RP0/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 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.

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/RP0/CPU0:router(config-rpl)# if destination in pset1 then
RP/0/RP0/CPU0:router(config-rpl-if)# drop
```



```
RP/0/RP0/CPU0:router(config-rpl-if)# endif  
RP/0/RP0/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 XR EXEC 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. |
| <i>name</i> | 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. |
| <i>set-element</i> | Value of the set element. |
| | <p>Note To inline edit multiple set elements separated with comma, use quotes to club the entries as a single argument. Example:</p> <pre>edit extcommunity-set rt rt_set inline add "4:4,5:4"</pre> |

Command Default

Default editor is GNU nano text editor

Command Modes

XR EXEC

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | 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.

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/RP0/CPU0:router# edit route-policy policy_A
-----
== MicroEMACS 3.8b () == repl_edit.139281 ==
  if destination in (2001::/8) then
    drop
  endif
end-policy
!

== MicroEMACS 3.8b () == repl_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/RP0/CPU0:router#edit route-policy policy_B
== MicroEMACS 3.8b () == repl_edit.141738
route-policy policy_B
  set metric-type type_1
  if destination in (2001::/8) then
    drop
  endif
end-policy
!
== MicroEMACS 3.8b () == repl_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 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.

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/RP0/CPU0:router(config)#policy-global
RP/0/RP0/CPU0:router(config-rp-gl)# glbpathtype 'ebgp'
RP/0/RP0/CPU0:router(config-rp-gl)# glbtag '100'
RP/0/RP0/CPU0:router(config-rp-gl)# end-global
```

Related Commands

| Command | Description |
|---------------------------------------------|---------------------------------------------|
| policy-global, on page 1275 | 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 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.

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/RP0/CPU0:router (config) #route-policy med-to-local-pref
RP/0/RP0/CPU0:router (config-rpl) #if med eq 150 then
RP/0/RP0/CPU0:router (config-rpl-if) # set local-preference 10
RP/0/RP0/CPU0:router (config-rpl-if) # elseif med eq 200 then
RP/0/RP0/CPU0:router (config-elseif) # set local-preference 60
RP/0/RP0/CPU0:router (config-elseif) # elseif med eq 250 then
RP/0/RP0/CPU0:router (config-elseif) # set local-preference 0

RP/0/RP0/CPU0:router (config-elseif) # endif
RP/0/RP0/CPU0:router (config-rpl) # end-policy
```

Related Commands

| Command | Description |
|--------------------------------------------------|--------------------------------------------------------------------|
| route-policy (RPL), on page 1289 | 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 XR Config 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 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.

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/RP0/CPU0:router(config)# as-path-set aset1
RP/0/RP0/CPU0:router(config-as)# ios-regex '_42$',
RP/0/RP0/CPU0:router(config-as)# ios-regex '_127$'
```

```
RP/0/RP0/CPU0:router(config-as)# end-set
RP/0/RP0/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/RP0/CPU0:router(config)# rd-set my_rd_set
```

```
RP/0/RP0/CPU0:router (config-rd) # 172.16.0.0/16:*,
RP/0/RP0/CPU0:router (config-rd) # 172.17.0.0/16:100,
RP/0/RP0/CPU0:router (config-rd) # 192:*,
RP/0/RP0/CPU0:router (config-rd) # 192:100
RP/0/RP0/CPU0:router (config-rd) # end-set
```


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 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.

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/RP0/CPU0:router(config)# route-policy extcommunity-is-empty-example
RP/0/RP0/CPU0:router(config-rpl)# if extcommunity rt is-empty then
RP/0/RP0/CPU0:router(config-rpl-if)# set local-preference 100

RP/0/RP0/CPU0:router(config-rpl-if)# endif
RP/0/RP0/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-name*inline-*extcommunity-set*parameter}

Syntax Description

| | |
|--------------------------------|------------------------------------------------------------------------------------------------------|
| <i>extcommunity-set-name</i> | Name of an RT extended community set. |
| <i>inline-extcommunity-set</i> | Inline RT extended community set. The inline extended community set must be enclosed in parentheses. |
| <i>parameter</i> | 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 5.0.0 | 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 **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/RP0/CPU0:router(config)# extcommunity-set rt my-extcommunity-set
RP/0/RP0/CPU0:router(config-ext)# 10:615,
RP/0/RP0/CPU0:router(config-ext)# 10:6150,
RP/0/RP0/CPU0:router(config-ext)# 15.15.15.15:15
RP/0/RP0/CPU0:router(config-ext)# end-set

RP/0/RP0/CPU0:router(config)# route-policy my-extcommunity-set-example($tag,$ip)
RP/0/RP0/CPU0:router(config-rpl)# if extcommunity rt matches-any my-extcommunity-set then
RP/0/RP0/CPU0:router(config-rpl-if)# set local-preference 100
RP/0/RP0/CPU0:router(config-rpl-if)# elseif extcommunity rt matches-any (10:20, 10:$tag)
then
RP/0/RP0/CPU0:router(config-rpl-elseif)# set local-preference 200
RP/0/RP0/CPU0:router(config-rpl-elseif)# elseif extcommunity rt matches-any ($ip:$tag) then
RP/0/RP0/CPU0:router(config-rpl-elseif)# set local-preference 300
RP/0/RP0/CPU0:router(config-rpl-elseif)# elseif extcommunity rt matches-any (2.3.4.5:$tag)
then
RP/0/RP0/CPU0:router(config-rpl-elseif)# set local-preference 400
RP/0/RP0/CPU0:router(config-rpl-elseif)# endif
RP/0/RP0/CPU0:router(config-rpl)# end-policy
```

Related Commands

| Command | Description |
|-------------------------------------------------------------|-----------------------------------------------------------|
| extcommunity rt matches-every, on page 1232 | 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-name*inline-*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

Release 5.0.0

Modification

This command was introduced.

Usage Guidelines

To use this command, you must be in a user group associated with a task group that 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/RP0/CPU0:router(config)# extcommunity-set rt my-extcommunity-set
RP/0/RP0/CPU0:router(config-ext)# 10:20,
RP/0/RP0/CPU0:router(config-ext)# 10:30,
RP/0/RP0/CPU0:router(config-ext)# 10:40
RP/0/RP0/CPU0:router(config-ext)# end-set
RP/0/RP0/CPU0:router(config)# route-policy extcommunity-matches-every-example($as,$tag)
RP/0/RP0/CPU0:router(config-rpl)# if extcommunity rt matches-every my-extcommunity-set then
RP/0/RP0/CPU0:router(config-rpl-if)# set local-preference 100
RP/0/RP0/CPU0:router(config-rpl-if)# elseif extcommunity rt matches-every (10:20, 10:$tag,
  $as:30) then
RP/0/RP0/CPU0:router(config-rpl-elseif)# set local-preference 200
RP/0/RP0/CPU0:router(config-rpl-elseif)# elseif
RP/0/RP0/CPU0:router(config-rpl-elseif)# set local-preference 300
RP/0/RP0/CPU0:router(config-rpl-elseif)# endif
RP/0/RP0/CPU0:router(config-rpl)# end-policy
```

Related Commands

| Command | Description |
|-----------------------------------------------------------|---------------------------------------------------------|
| extcommunity rt matches-any, on page 1230 | 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-nameinline-extcommunity-setparameter}*

| Syntax Description | |
|--------------------------------------|------------------------------------------------------------|
| <i>rt-type-extcommunity-set-name</i> | Name of an RT extended community set. |
| <i>inline-extcommunity-set</i> | Inline RT extended community set, enclosed in parentheses. |
| <i>parameter</i> | Parameter name preceded with a "\$" symbol. |

Command Default None

Command Modes Route-policy configuration

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | 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.

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.

```
RP/0/RP0/CPU0:router(config)#extcommunity-set rt my-extcommunity-set
RP/0/RP0/CPU0:router(config-ext)#10:615,
RP/0/RP0/CPU0:router(config-ext)#10:6150,
RP/0/RP0/CPU0:router(config-ext)#15.15.15.15:15
RP/0/RP0/CPU0:router(config-ext)#end-set
RP/0/RP0/CPU0:router(config)#route-policy my-extcommunity-set-example($tag,$ip)
RP/0/RP0/CPU0:router(config-rpl)#if extcommunity rt matches-within my-extcommunity-set then
RP/0/RP0/CPU0:router(config-rpl-if)#set local-preference 100
```

extcommunity-set cost

To define a cost extended community set, use the **extcommunity-set cost** command in XR Config 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

XR Config

Command History

| Release | Modification |
|---------------|------------------------------|
| 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.

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:

- **#-remark** ---Remark beginning with '#'
- **0-255**---Decimal number
- **abort** ---Discard RPL definition and return to top level config
- **end-set** ---End of set definition
- **exit** ---Exit from the submode
- **igp**---Cost Community with IGP as point of insertion
- **pre-bestpath**: ---Cost Community with Pre-Bestpath as point of insertion
- **show** ---Show partial RPL configuration

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/RP0/CPU0:router(config)# extcommunity-set cost extcomm-cost  
RP/0/RP0/CPU0:router(config-ext)# IGP:90:914,  
RP/0/RP0/CPU0:router(config-ext)# Pre-Bestpath:91:915  
RP/0/RP0/CPU0:router(config-ext)# end-set
```

extcommunity-set rt

To define a Border Gateway Protocol (BGP) route target (RT) extended community set, use the **extcommunity-set rt** command in XR Config 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 | <i>name</i> Name of an RT extended community set. | | | | |
|---------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------|--------------|---------------|------------------------------|
| Command Default | No default behavior or values | | | | |
| Command Modes | XR Config | | | | |
| Command History | <table border="1"> <thead> <tr> <th>Release</th> <th>Modification</th> </tr> </thead> <tbody> <tr> <td>Release 5.0.0</td> <td>This command was introduced.</td> </tr> </tbody> </table> | Release | Modification | Release 5.0.0 | This command was introduced. |
| Release | Modification | | | | |
| 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.

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:

- *#-remark* ---Remark beginning with '#'
- ***--- 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



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/RP0/CPU0:router(config)# extcommunity-set rt extcomm-rt
RP/0/RP0/CPU0:router(config-ext)# 10002:666
RP/0/RP0/CPU0:router(config-ext)# 10.0.0.2:666
RP/0/RP0/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 XR Config 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 | <i>name</i> Name of an SoO extended community set. |
|---------------------------|----------------------------------------------------|

| | |
|------------------------|-------------------------------|
| Command Default | No default behavior or values |
|------------------------|-------------------------------|

| | |
|----------------------|-----------|
| Command Modes | XR Config |
|----------------------|-----------|

| Command History | Release | Modification |
|------------------------|----------------|------------------------------|
| | 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. |
|-------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

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:

- *#-remark* ---Remark beginning with '#'
- ***--- 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
- **abort** ---Discard RPL definition and return to top level config
- **dfa-regex** ---DFA style regular expression
- **end-set** ---End of set definition
- **exit** ---Exit from the submode
- **ios-regex** ---Traditional IOS style regular expression

- **show** ---Show partial RPL configuration

N is a site-specific number.

Examples

In the following example, a SoO extended community set named extcomm-soo is defined:

```
RP/0/RP0/CPU0:router(config)# extcommunity-set soo extcomm-soo  
RP/0/RP0/CPU0:router(config-ext)# 66:60001,  
RP/0/RP0/CPU0:router(config-ext)# 77:70001,  
RP/0/RP0/CPU0:router(config-ext)# 88:80001,  
RP/0/RP0/CPU0:router(config-ext)# 99:90001,  
  
RP/0/RP0/CPU0:router(config-ext)# 100.100.100.1:153  
RP/0/RP0/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 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.

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/RP0/CPU0:router(config)# route-policy extcommunity-is-empty-example
RP/0/RP0/CPU0:router(config-rpl)# if extcommunity soo is-empty then
RP/0/RP0/CPU0:router(config-rpl-if)# set local-preference 100
RP/0/RP0/CPU0:router(config-rpl-if)# endif
RP/0/RP0/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-name*inline-*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

Release 5.0.0

Modification

This command was introduced.

Usage Guidelines

To use this command, you must be in a user group associated with a task group that 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/RP0/CPU0:router(config)# extcommunity-set soo extcomm-soo
RP/0/RP0/CPU0:router(config-ext)# 66:60001,
RP/0/RP0/CPU0:router(config-ext)# 77:70001,
RP/0/RP0/CPU0:router(config-ext)# 88:80001,
RP/0/RP0/CPU0:router(config-ext)# 99:90001,
RP/0/RP0/CPU0:router(config-ext)# 100.100.100.1:153
RP/0/RP0/CPU0:router(config-ext)# end-set

RP/0/RP0/CPU0:router(config)# route-policy my-extcommunity-set-example($tag,$ip)
RP/0/RP0/CPU0:router(config-rpl)# if extcommunity soo matches-any extcomm-soo then
RP/0/RP0/CPU0:router(config-rpl-if)# set local-preference 100
RP/0/RP0/CPU0:router(config-rpl-if)# elseif extcommunity soo matches-any (10:20, 10:$tag)
then
RP/0/RP0/CPU0:router(config-rpl-elseif)# set local-preference 200
RP/0/RP0/CPU0:router(config-rpl-elseif)# elseif extcommunity soo matches-any ($ip:$tag)
then
RP/0/RP0/CPU0:router(config-rpl-elseif)# set local-preference 300
RP/0/RP0/CPU0:router(config-rpl-elseif)# elseif extcommunity soo matches-any (2.3.4.5:$tag)
then
RP/0/RP0/CPU0:router(config-rpl-elseif)# set local-preference 400
RP/0/RP0/CPU0:router(config-rpl-elseif)# endif
RP/0/RP0/CPU0:router(config-rpl)# end-policy
```

Related Commands

| Command | Description |
|--------------------------------------------------------------|------------------------------------------------------------|
| extcommunity rt matches-any, on page 1230 | Matches any element of a BGP RT extended community set. |
| extcommunity soo matches-every, on page 1245 | 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-name*inline-*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

Release 5.0.0

Modification

This command was introduced.

Usage Guidelines

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

extcommunity soo matches-every

```

RP/0/RP0/CPU0:router(config)# extcommunity-set soo my-extcomm-rt-set
RP/0/RP0/CPU0:router(config-ext)# 10:20,
RP/0/RP0/CPU0:router(config-ext)# 10:30,
RP/0/RP0/CPU0:router(config-ext)# 10:40
RP/0/RP0/CPU0:router(config-ext)# end-set

RP/0/RP0/CPU0:router(config)# route-policy extcommunity-matches-every-example($as, $tag)
RP/0/RP0/CPU0:router(config-rpl)# if extcommunity soo matches-every my-extcomm-rt-set then
RP/0/RP0/CPU0:router(config-rpl-if)# set local-preference 100
RP/0/RP0/CPU0:router(config-rpl-if)# elseif extcommunity soo matches-every (10:20, 10:$tag,
$as:30) then
RP/0/RP0/CPU0:router(config-rpl-elseif)# set local-preference 200
RP/0/RP0/CPU0:router(config-rpl-elseif)# else
RP/0/RP0/CPU0:router(config-rpl-elseif)# set local-preference 300
RP/0/RP0/CPU0:router(config-rpl-elseif)# endif
RP/0/RP0/CPU0:router(config-rpl)# end-policy

```

Related Commands

| Command | Description |
|------------------------------------------------------------|----------------------------------------------------------|
| extcommunity soo matches-any, on page 1243 | Matches any element of a BGP SoO extended community set. |

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

| | |
|-------------------------------|-----------------------------------------------------------------------------------------|
| <i>conditional-expression</i> | 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. |
| <i>action-statement</i> | Sequence of operations that modify a route. |

Command Default

No default behavior or values

Command Modes

Route-policy configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| 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.

The **if** command uses a conditional expression to decide which actions or dispositions should be taken for a given route. [Table 143: Conditional Expressions, on page 1248](#) 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 144: Action Statements, on page 1249](#) lists the action statements.

Apply Condition policies allow usage of a route-policy in an "if" statement of another route-policy.

```
Route-policy policy_name
If apply policyA and apply policyB then
Set med 100
Else if not apply policyD then
Set med 200
Else
Set med 300
Endif
End-policy
```

Table 143: Conditional Expressions

| Command | Description |
|--------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------|
| as-path in, on page 1183 | 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 1185 | Determines if the router (or another router within this autonomous system or confederation) originated the route. |
| as-path length, on page 1187 | Performs a conditional check based on the length of the AS path. |
| as-path neighbor-is, on page 1189 | 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 1192 | Tests an AS path against the AS sequence beginning with the AS number that originated a route. |
| as-path passes-through, on page 1194 | 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 1198 | Performs specific checks based on the length of the AS path. |
| community is-empty, on page 1200 | Learns if a route has community attributes associated with it. |
| community matches-any, on page 1201 | Matches any element of a community set. |
| community matches-every, on page 1203 | Matches every element of a community set. |
| destination in, on page 1216 | Matches a destination entry in a named prefix set or inline prefix set. |
| extcommunity rt is-empty, on page 1229 | Learns if a route has RT extended community attributes associated with it. |
| extcommunity rt matches-any, on page 1230 | Matches elements of an RT extended community set. |
| extcommunity rt matches-every, on page 1232 | Matches every element of an RT extended community set. |
| extcommunity rt matches-within, on page 1234 | Matches at least one element of a Border Gateway Protocol (BGP) route target (RT) extended community set. |
| extcommunity soo is-empty, on page 1242 | Learns if a route has SoO extended community attributes associated with it. |
| extcommunity soo matches-any, on page 1243 | Matches elements of an SoO extended community set. |
| extcommunity soo matches-every, on page 1245 | Matches every element of an SoO extended community set. |

| Command | Description |
|------------------------------------------------|--------------------------------------------------------------------------------------------------------------|
| local-preference, on page 1253 | Specifies BGP local-preference attribute |
| med, on page 1263 | Compares the MED to an integer value or a parameterized value. |
| next-hop in, on page 1264 | 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 1266 | Matches a prefix in a prefix set or an inline prefix set. |
| origin is, on page 1268 | Tests the value of the origin attribute. |
| path-type is, on page 1274 | Tests the path type. |
| protocol, on page 1281 | Checks if a protocol is installing the route. |
| rd in, on page 1283 | Compares the RD associated with the route to data contained in either a named or an inline RD set. |
| rib-has-route, on page 1286 | Checks if a route is in the RIB. |
| route-has-label, on page 1288 | Checks if a route has a Multiprotocol Label Switching (MPLS) label. |
| route-type is, on page 1291 | Compares route types when redistribution is being performed into BGP, OSPF, or IS-IS. |
| source in, on page 1432 | Tests the source of the route against the data in either a named or an inline prefix set. |
| tag, on page 1436 | Matches a specific tag value. |
| tag in, on page 1437 | Conditionally compares tag-route against tag-set. |

Table 144: Action Statements

| Command | Description |
|------------------------------------------------------|-----------------------------------------------------------------------------------|
| abort (RPL), on page 1177 | Discards a route policy definition and returns to XR Config mode. |
| add, on page 1179 | Adds an offset to an existing value. |
| apply, on page 1181 | Executes a parameterized or an unparameterized policy from within another policy. |
| delete community, on page 1210 | Deletes community values from a community list in a route. |
| delete extcommunity rt, on page 1214 | Deletes extended community values from an extended community list in a route. |
| done, on page 1218 | Accepts this route with no further processing |
| drop, on page 1220 | Drops a route. |

| Command | Description |
|-------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------|
| end-policy, on page 1226 | Ends the definition of a route policy and exits route-policy configuration mode. |
| pass, on page 1272 | 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 1279 | Prepends the AS path with additional autonomous system numbers. |
| replace as-path, on page 1284 | Replaces a sequence of AS numbers or private AS numbers in the AS path with the configured local AS. |
| set community, on page 1297 | Sets the BGP community attribute. |
| set dampening, on page 1299 | Configures BGP route dampening. |
| set eigrp-metric, on page 1301 | Sets the Enhanced Interior Gateway Routing Protocol (EIGRP) metric value. |
| set extcommunity cost, on page 1303 | Replaces or adds the extended communities for a cost on the route. |
| set extcommunity rt, on page 1305 | Replaces or adds the extended communities for an RT on the route. |
| set ip-precedence, on page 1307 | Sets the IP precedence to classify packets. |
| set isis-metric, on page 1309 | Sets the IS-IS metric attribute value. |
| set label, on page 1310 | Sets the BGP label attribute value. |
| set level, on page 1313 | Configures the IS-IS level in which redistributed routes should be sent. |
| set local-preference, on page 1314 | Specifies a preference value for the autonomous system path. |
| set med, on page 1315 | Sets the MED value. |
| set metric-type (IS-IS), on page 1317 | Controls whether IS-IS treats the metric as an internal or external metric. |
| set metric-type (OSPF), on page 1318 | Controls whether OSPF treats the cost as a Type 1 or Type 2 metric. |
| set next-hop, on page 1319 | Replaces the next-hop associated with a given route. |
| set origin, on page 1321 | Changes the origin attribute. |
| set ospf-metric, on page 1322 | Sets an OSPF protocol metric attribute value. |
| set qos-group (RPL), on page 1325 | Sets the QoS group to classify packets. |
| set rib-metric, on page 1326 | Sets a RIB metric attribute value for a table policy. |
| set rip-metric, on page 1327 | Sets RIP metric attributes. |
| set rip-tag, on page 1328 | Sets route tag attribute. |
| set tag, on page 1331 | Sets the tag attribute. |

| Command | Description |
|-------------------------------------------------|-------------------------------------------------------------------------------------------------|
| set traffic-index, on page 1332 | Sets the traffic index attribute. |
| set weight, on page 1334 | Sets the weight value for BGP routes. |
| suppress-route, on page 1435 | Indicates that a given component of an aggregate should be suppressed, that is, not advertised. |
| unsuppress-route, on page 1440 | Indicates that a given component of an aggregate should be unsuppressed. |

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/RP0/CPU0:router(config-rpl)# if as-path in as-path-set-1 then
RP/0/RP0/CPU0:router(config-rpl-if)# drop
RP/0/RP0/CPU0:router(config-rpl-if)# endif
RP/0/RP0/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/RP0/CPU0:router(config-rpl)# if origin is igp then
RP/0/RP0/CPU0:router(config-rpl-if)# set med 42
RP/0/RP0/CPU0:router(config-rpl-if)# prepend as-path 73 5
RP/0/RP0/CPU0:router(config-rpl-if)# endif
RP/0/RP0/CPU0:router(config-rpl)#
```

The **if** command also permits an **else** clause to be executed if the expression is false, as follows:

```
RP/0/RP0/CPU0:router(config-rpl)# if med eq 200 then
RP/0/RP0/CPU0:router(config-rpl-if)# set community (12:34) additive
RP/0/RP0/CPU0:router(config-rpl-if)# else
RP/0/RP0/CPU0:router(config-rpl-else)# set community (12:56) additive
RP/0/RP0/CPU0:router(config-rpl-else)# endif
RP/0/RP0/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/RP0/CPU0:router(config-rpl)# if med eq 150 then
RP/0/RP0/CPU0:router(config-rpl-if)# set local-preference 10
RP/0/RP0/CPU0:router(config-rpl-if)# elseif med eq 200 then
RP/0/RP0/CPU0:router(config-rpl-elseif)# set local-preference 60
RP/0/RP0/CPU0:router(config-rpl-elseif)# elseif med eq 250 then
RP/0/RP0/CPU0:router(config-rpl-elseif)# set local-preference 110
```

```
RP/0/RP0/CPU0:router(config-rpl-elseif)# else
RP/0/RP0/CPU0:router(config-rpl-else)# set local-preference 0
RP/0/RP0/CPU0:router(config-rpl-else)# endif
RP/0/RP0/CPU0:router(config-rpl)#
```

The statements within an **if** statement may themselves be **if** statements, as shown in this example:

```
RP/0/RP0/CPU0:router(config-rpl)# if community matches-any (12:34, 56:78) then
RP/0/RP0/CPU0:router(config-rpl-if)# if med eq 150 then
RP/0/RP0/CPU0:router(config-rpl-if)# drop
RP/0/RP0/CPU0:router(config-rpl-if)# endif
RP/0/RP0/CPU0:router(config-rpl-if)# set local-preference 100
RP/0/RP0/CPU0:router(config-rpl-if)# endif
RP/0/RP0/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 Discriminator (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.

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. |
| <i>number</i> | Value assigned to a 32-bit unsigned integer. Range is 0 to 4294967295. |
| <i>parameter</i> | 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 5.0.0 | 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 **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/RP0/CPU0:router#config
RP/0/RP0/CPU0:router (config)#route-policy lrg_comm_rp4
RP/0/RP0/CPU0:router (config-rpl)#if large-community is-empty then
RP/0/RP0/CPU0:router (config-rpl)#set local-preference 104
```

```
RP/0/RP0/CPU0:router(config-rpl)#endif  
RP/0/RP0/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 neighbor 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/RP0/CPU0:router#config
RP/0/RP0/CPU0:router (config)#route-policy elbonia
RP/0/RP0/CPU0:router (config-rpl)#if large-community matches-any (1:2:3, 4:5:*) then
RP/0/RP0/CPU0:router (config-rpl)#set local-preference 94
RP/0/RP0/CPU0:router (config-rpl)#endif
RP/0/RP0/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 neighbor 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/RP0/CPU0:router#config
RP/0/RP0/CPU0:router(config)#route-policy bob
RP/0/RP0/CPU0:router(config-rpl)#if large-community matches-any (*:.*:3, 4:5:*) then
RP/0/RP0/CPU0:router(config-rpl)#set local-preference 94
RP/0/RP0/CPU0:router(config-rpl)#endif
RP/0/RP0/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

| | |
|-----------------------------------|---------------------------------------------------------------------------------------------|
| <i>large-community-set-name</i> | Name of a large community set. |
| <i>inline-large-community-set</i> | Inline large community set. The inline large community set must be enclosed in parentheses. |
| <i>parameter</i> | 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 neighbor 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/RP0/CPU0:router#config
RP/0/RP0/CPU0:router(config)#route-policy bob
RP/0/RP0/CPU0:router(config-rpl)#if large-community matches-within (*:*:3, 4:5:*) then
RP/0/RP0/CPU0:router(config-rpl)#set local-preference 103
RP/0/RP0/CPU0:router(config-rpl)#endif
RP/0/RP0/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 XR Config 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

XR Config

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/RP0/CPU0:router#configure
RP/0/RP0/CPU0:router(config)#large-community-set catbert
RP/0/RP0/CPU0:router(config-largecomm)#1:2:3,
RP/0/RP0/CPU0:router(config-largecomm)#[5..9]:2:3
RP/0/RP0/CPU0:router(config-largecomm)#1:3:*
RP/0/RP0/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. |
| <i>number</i> | Value assigned to a 32-bit unsigned integer. Range is 0 to 4294967295. |
| <i>parameter</i> | 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 5.0.0 | 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 **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/RP0/CPU0:router(config-rpl)# if med eq 10 then  
RP/0/RP0/CPU0:router(config-rpl-if)# set local-preference 100  
RP/0/RP0/CPU0:router(config-rpl-if)# endif  
RP/0/RP0/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-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 name. The parameter name must be preceded with a "\$."

Command Default

No default behavior or values

Command Modes

Route-policy configuration

Command History

Release

Release 5.0.0

Modification

This command was introduced.

Usage Guidelines

To use this command, you must be in a user group associated with a task group that 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/RP0/CPU0:router(config-rpl)# if next-hop in some-prefix-set then
RP/0/RP0/CPU0:router(config-rpl-if)# if next-hop in (10.0.0.5, fe80::230/64) then
RP/0/RP0/CPU0:router(config-rpl-if)# set local-preference 0
```

```
RP/0/RP0/CPU0:router(config-rpl-if)# endif  
RP/0/RP0/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 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.

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/RP0/CPU0:router (config) # prefix-set orfpreset1
RP/0/RP0/CPU0:router (config-pfx) # 211.105.1.0/24,
RP/0/RP0/CPU0:router (config-pfx) # 211.105.5.0/24,
RP/0/RP0/CPU0:router (config-pfx) # 211.105.11.0/24
RP/0/RP0/CPU0:router (config-pfx) # end-set
!
!
RP/0/RP0/CPU0:router (config) # route-policy orfpolicy
```

```

RP/0/RP0/CPU0:router(config-rpl)# if orf prefix in orfpreset1 then
RP/0/RP0/CPU0:router(config-rpl-if)# drop
RP/0/RP0/CPU0:router(config-rpl-if)# endif
RP/0/RP0/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/RP0/CPU0:router(config-rpl-if)# pass
RP/0/RP0/CPU0:router(config-rpl-if)# endif
RP/0/RP0/CPU0:router(config-rpl)# end-policy
!
!
RP/0/RP0/CPU0:router(config)# router bgp 2
RP/0/RP0/CPU0:router(config-bgp)# neighbor 1.1.1.1
RP/0/RP0/CPU0:router(config-bgp-nbr)# remote-as 3
RP/0/RP0/CPU0:router(config-bgp-nbr)# address-family ipv4 unicast
RP/0/RP0/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. |
| <i>parameter</i> | 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 5.0.0 | 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 **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/RP0/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/RP0/CPU0:router(config)# route-policy bar($origin)
RP/0/RP0/CPU0:router(config-rpl)# if origin is $origin then
RP/0/RP0/CPU0:router(config-rpl-if)# set med 20
RP/0/RP0/CPU0:router(config-rpl-if)# endif
RP/0/RP0/CPU0:router(config-rpl)#
```

ospf-area-set

Defines an OSPF area set to be used in routing policy statements.

```
ospf-area-setname {<0-4294967295><A.B.C.D> | abort | end-set | exit | show}
noospf-area-setname
```

| 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

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | Release 5.0.0 | 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/RP0/CPU0:router (config) # ospf-area-set ospf_area_set_demo1
RP/0/RP0/CPU0:router (config-ospf-area) # 10.0.0.1,
RP/0/RP0/CPU0:router (config-ospf-area) # 3553
RP/0/RP0/CPU0:router (config-ospf-area) # end-set

RP/0/RP0/CPU0:router (config) # ospf-area-set ospf_area_set_demo2
RP/0/RP0/CPU0:router (config-ospf-area) # 20.0.0.2,
RP/0/RP0/CPU0:router (config-ospf-area) # 3673
RP/0/RP0/CPU0:router (config-ospf-area) # end-set

RP/0/RP0/CPU0:router (config) # route-policy use_ospf_area_set
RP/0/RP0/CPU0:router (config-rpl) # if ospf-area in ospf-area-set* then set ospf-metric 200
RP/0/RP0/CPU0:router (config-rpl-if) # elseif ospf-area in( 10.0.0.1, 10.0.0.2 )then set
ospf-metric 300
```

```
RP/0/RP0/CPU0:router(config-rpl-elseif)# endif  
RP/0/RP0/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 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.

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/RP0/CPU0:router(config-rpl)# pass
```

This example accepts the route unconditionally, without modifying it, if the destination is in prefix-set permitted:

```
RP/0/RP0/CPU0:router(config-rpl)# if destination in permitted then
RP/0/RP0/CPU0:router(config-rpl-if)# pass
RP/0/RP0/CPU0:router(config-rpl-if)# endif
```

```
RP/0/RP0/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. |
| | <i>parameter</i> | 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 5.0.0 | 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 **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/RP0/CPU0:router(config)# route-policy policy_A
RP/0/RP0/CPU0:router(config-rpl)# if path-type is ebgp then
RP/0/RP0/CPU0:router(config-rpl-if)# pass
RP/0/RP0/CPU0:router(config-rpl-if)# else
RP/0/RP0/CPU0:router(config-rpl-else)# drop
RP/0/RP0/CPU0:router(config-rpl-if)# endif
RP/0/RP0/CPU0:router(config-rpl)# end-policy
```

policy-global

To define global parameters and enter global parameter configuration mode, use the **policy-global** command in XR Config 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 XR Config

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | 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.

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 NCS 6000 Series Routers*

| Task ID | Task ID | Operations |
|---------|--------------|----------------|
| | route-policy | read, write |

Examples

The following example shows how to configure global parameters:

```
RP/0/RP0/CPU0:router(config)# policy-global
RP/0/RP0/CPU0:router(config-rp-gl)# glbpathtype 'ebgp'
RP/0/RP0/CPU0:router(config-rp-gl)# glbtag '100'
RP/0/RP0/CPU0:router(config-rp-gl)# end-global
```

In the following example, the *globalparam* argument makes use of the global parameters *glbpathtype* and *glbtag* defined above and is defined for a nonparameterized policy:

```
RP/0/RP0/CPU0:router(config)# route-policy globalparam
RP/0/RP0/CPU0:router(config-rpl)# if path-type is $glbpathtype then
```

```
RP/0/RP0/CPU0:router(config-rpl)# set tag $glbtag  
RP/0/RP0/CPU0:router(config-rpl)# endif  
RP/0/RP0/CPU0:router(config-rpl)# end-policy
```

Related Commands

| Command | Description |
|------------------------------------------|-------------------------------------------|
| end-global, on page 1225 | 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 XR Config mode. To remove a named prefix set, use the **no** form of this command.

```
prefix-set name
no prefix-set name
```

| | | |
|---------------------------|----------------|------------------------------|
| Syntax Description | <i>name</i> | Name of a prefix set. |
| Command Default | None | |
| Command Modes | XR Config | |
| Command History | Release | Modification |
| | 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.

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 **greater than or equal 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 **less than or equal 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 **equal 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/RP0/CPU0:router (config) # prefix-set legal-ipv4-prefix-examples
RP/0/RP0/CPU0:router (config-pfx) # 10.0.1.1,
RP/0/RP0/CPU0:router (config-pfx) # 10.0.2.0/24,
RP/0/RP0/CPU0:router (config-pfx) # 10.0.3.0/24 ge 28,
RP/0/RP0/CPU0:router (config-pfx) # 10.0.4.0/24 le 28,
RP/0/RP0/CPU0:router (config-pfx) # 10.0.5.0/24 ge 26 le 30,
RP/0/RP0/CPU0:router (config-pfx) # 10.0.6.0/24 eq 28
RP/0/RP0/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/RP0/CPU0:router (config) # prefix-set INVALID-PREFIX-EXAMPLES
RP/0/RP0/CPU0:router (config-pfx) # 10.1.1.1 ge 16,
RP/0/RP0/CPU0:router (config-pfx) # 10.1.2.1 le 16,
RP/0/RP0/CPU0:router (config-pfx) # 10.1.3.0/24 le 23,
RP/0/RP0/CPU0:router (config-pfx) # 10.1.4.0/24 ge 33,
RP/0/RP0/CPU0:router (config-pfx) # 10.1.5.0/25 ge 29 le 28
RP/0/RP0/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/RP0/CPU0:router (config) # prefix-set legal-ipv6-prefix-examples
RP/0/RP0/CPU0:router (config-pfx) # 2001:0:0:1::/64,
RP/0/RP0/CPU0:router (config-pfx) # 2001:0:0:2::/64,
RP/0/RP0/CPU0:router (config-pfx) # 2001:0:0:3::/64,
RP/0/RP0/CPU0:router (config-pfx) # 2001:0:0:4::/64
RP/0/RP0/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

| | |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>as-number</i> | Autonomous system number to prepend to the path. <ul style="list-style-type: none"> 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. |
| <i>parameter</i> | Parameter name. The parameter name must be preceded with a "\$." |
| most-recent | Specifies that the most recent autonomous system number should be prepended. |
| <i>number</i> | (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 5.0.0 | This command was introduced. |

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/RP0/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/RP0/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 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.

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/RP0/CPU0:router(config)# route-policy rip1
RP/0/RP0/CPU0:router(config-rpl)# if protocol in (connected, static) then
RP/0/RP0/CPU0:router(config-rpl-if)# add rip-metric 2
RP/0/RP0/CPU0:router(config-rpl-if)# elseif protocol is bgp 1 then
RP/0/RP0/CPU0:router(config-rpl-elseif)# add rip-metric 3
RP/0/RP0/CPU0:router(config-rpl-elseif)# elseif protocol is ospf 2 then
RP/0/RP0/CPU0:router(config-rpl-elseif)# add rip-metric 4
RP/0/RP0/CPU0:router(config-rpl-elseif)# else
RP/0/RP0/CPU0:router(config-rpl-else)# add rip-metric 5
RP/0/RP0/CPU0:router(config-rpl-else)# endif
RP/0/RP0/CPU0:router(config-rpl)# end-policy

RP/0/RP0/CPU0:router(config)# router rip
RP/0/RP0/CPU0:router(config-rip)# interface GigabitEthernet0/1/0/1
RP/0/RP0/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-name|inline-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 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.

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/RP0/CPU0:router(config)# route-policy
RP/0/RP0/CPU0:router(config-rpl)# if rd in (128.1.0.0/16:100) then
RP/0/RP0/CPU0:router(config-rpl-if)# pass
RP/0/RP0/CPU0:router(config-rpl-if)# endif
RP/0/RP0/CPU0:router(config-rpl)# end-policy
```

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 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.

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/RP0/CPU0:router(config)# route-policy drop-as-1234
RP/0/RP0/CPU0:router(config-rpl)# replace as-path '90 78 45 $asnum'
RP/0/RP0/CPU0:router(config-rpl)# replace as-path private-as
RP/0/RP0/CPU0:router(config-rpl)# replace as-path '9.9 7.89 14.15 $asnum'
RP/0/RP0/CPU0:router(config-rpl)# replace as-path '9 89 14.15 $asnum'
```

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 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 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 5.0.0 | 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 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 le 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 is used to learn if a route contained in a prefix set 10.10.0.0/16 is in the RIB:

```
RP/0/RP0/CPU0:router(config-rpl)# if rib-has-route in (10.10.0.0/16 ge 16) then
RP/0/RP0/CPU0:router(config-rpl-if)# pass
RP/0/RP0/CPU0:router(config-rpl-if)# endif
```

```
RP/0/RP0/CPU0:router(config-rpl)#
```

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 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.

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/RP0/CPU0:router(config-rpl)# if route-has-label then
RP/0/RP0/CPU0:router(config-rpl-if)# pass
RP/0/RP0/CPU0:router(config-rpl-if)# endif
RP/0/RP0/CPU0:router(config-rpl)#
```

route-policy (RPL)

To define a route policy and enter route-policy configuration mode, use the **route-policy** command in XR Config 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

XR Config

Command History

Release

Release 5.0.0

Modification

This command was introduced.

Usage Guidelines

To use this command, you must be in a user group associated with a task group that 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/RP0/CPU0:router(config)# route-policy drop-everything
RP/0/RP0/CPU0:router(config-rpl)# drop
RP/0/RP0/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/RP0/CPU0:router(config)# route-policy drop-as-1234
RP/0/RP0/CPU0:router(config-rpl)# if as-path passes-through '1234' then
RP/0/RP0/CPU0:router(config-rpl-if)# apply check-communities
RP/0/RP0/CPU0:router(config-rpl-if)# else
RP/0/RP0/CPU0:router(config-rpl-else)# pass
RP/0/RP0/CPU0:router(config-rpl-else)# endif
RP/0/RP0/CPU0:router(config-rpl)# end-policy
```

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 1226 | 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-1** | **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. |
| <i>parameter</i> | 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 5.0.0 | 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 **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/RP0/CPU0:router(config)# route-policy policy_A
RP/0/RP0/CPU0:router(config-rpl)# if route-type is local then
RP/0/RP0/CPU0:router(config-rpl-if)# pass
RP/0/RP0/CPU0:router(config-rpl-if)# else
RP/0/RP0/CPU0:router(config-rpl-else)# drop
RP/0/RP0/CPU0:router(config-rpl-if)# endif
RP/0/RP0/CPU0:router(config-rpl)# end-policy
```


rpl editor

To set the default routing policy language (RPL) editor, use the **rpl editor** command in XR Config 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

XR Config

Command History

Release

Modification

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 Operations

route-policy read,
write

Examples

In the following example, the default RPL editor is set to Nano:

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# rpl editor nano
```

In the following example, the default RPL editor is set to EMACS:

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# rpl editor emacs
```

In the following example, the default RPL editor is set to VIM:

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# rpl editor vim
```

rpl maximum

To configure system limits on the routing policy subsystem, use the **rpl maximum** command in XR Config mode.

rpl maximum {**lines** | **policies**} *number*

| Syntax Description | lines <i>number</i> | Configures the number of lines of configuration limit. Range is from 1 to 131072. |
|--------------------|-------------------------------|-----------------------------------------------------------------------------------|
| | policies <i>number</i> | Configures the number of policies limit. Range is from 1 to 5000. |

| Command Default | lines <i>number</i> : 65536 |
|-----------------|---------------------------------------|
| | policies <i>numbers</i> : 3500 |

| Command Modes | XR Config |
|---------------|-----------|
|---------------|-----------|

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | 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.

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/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# rpl maximum lines 50
RP/0/RP0/CPU0:router(config)# rpl maximum policies 6
```

Related Commands

| Command | Description |
|------------------------------------------------|--------------------------------------------------------------------------------|
| show rpl maximum, on page 1388 | Displays the maximum limits for lines of configuration and number of policies. |

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 aigp-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 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-policy | read, write |

Examples The following example shows how to set the aigp metric as the igp cost for route-policy aigp_policy:

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# route-policy aigp_policy
RP/0/RP0/CPU0:router(config-rpl)# set aigp-metric igp-cost
```

| Related Commands | Command | 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-name**inline-community-set**parameter*} [**additive**]

Syntax Description

| | |
|-----------------------------|---------------------------------------------------------------------------------|
| <i>community-set-name</i> | Community set name. |
| <i>inline-community-set</i> | Inline community set. The inline community set must be enclosed in parentheses. |
| <i>parameter</i> | 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 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.

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 **peer-as** 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 **peer-as** 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/RP0/CPU0:router(config-rpl)# set community (10:24)
RP/0/RP0/CPU0:router(config-rpl)# set community (10:24, $as:24, $as:$tag)
RP/0/RP0/CPU0:router(config-rpl)# set community (10:24, internet) additive
RP/0/RP0/CPU0:router(config-rpl)# set community (10:24, $as:24) additive
RP/0/RP0/CPU0:router(config-rpl)# set community (10:24, peeras:24) additive
```

set dampening

To configure Border Gateway Protocol (BGP) route dampening, use the **set dampening** command in route-policy configuration mode.

```
set dampening {half-life {minutesparameter} | max-suppress {minutesparameter} | reuse {secondsparameter} | suppress {penalty-unitsparameter} | others default}
```

| Syntax Description | Parameter | Description |
|--------------------------------------|------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| half-life <i>minutes</i> | | 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. |
| | <i>parameter</i> | Parameter name. The parameter name must be preceded with a "\$." |
| max-suppress <i>minutes</i> | | 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 <i>seconds</i> | | 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 <i>penalty-units</i> | | 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 <i>must</i> end with others default . This designation indicates that any keyword not defined is set to its default. |

| Command Default | Value |
|---------------------|---------------------------------------|
| half-life | 15 minutes |
| max-suppress | 60 minutes (four times the half-life) |
| reuse | 750 seconds |
| suppress | 2000 penalty units |

| Command Modes | Mode |
|---------------|----------------------------|
| | Route-policy configuration |

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | 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.

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/RP0/CPU0:router(config-rpl)# set dampening halflife 20 others default
RP/0/RP0/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/RP0/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/RP0/CPU0:router(config-rpl)# set dampening reuse 700
```

In the following example, the parameters are used.

```
RP/0/RP0/CPU0:router(config-rpl)# set dampening halflife $p1 suppress $p4 reuse $p3
max-suppress $p2
```


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 | |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>bandwidth</i> | Minimum bandwidth of the route in kilobits per second. Range is 1 to 4294967295. |
| <i>delay</i> | 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. |
| <i>reliability</i> | Likelihood of successful packet transmission expressed as a number between 0 and 255. The value 255 means 100 percent reliability; 0 means no reliability. |
| <i>loading</i> | Effective bandwidth of the route expressed as a number from 1 to 255 (255 is 100 percent loading). |
| <i>mtu</i> | 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 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.

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/RP0/CPU0:router(config)# route-policy policy_1
RP/0/RP0/CPU0:router(config-rpl)# set eigrp-metric 1400 120 250 100 1500
RP/0/RP0/CPU0:router(config-rpl)# end-policy
```

Related Commands

| Command | Description |
|-----------------------------------|--------------------------------------|
| add, on page 1179 | Adds an offset to an existing value. |

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

| | |
|-------------------------------------|-------------------------------------------------------------------------------------------------------------|
| <i>cost-extcommunity-set-name</i> | Cost extended community set name. |
| <i>cost-inline-extcommunity-set</i> | Inline cost extended community set. The inline cost extended community set must be enclosed in parentheses. |
| <i>parameter</i> | 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 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.

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/RP0/CPU0:router(config-rpl)# set extcommunity cost (IGP:10:20)
RP/0/RP0/CPU0:router(config-rpl)# set extcommunity cost (Pre-Bestpath:33:44)
RP/0/RP0/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-name|rt-inline-extcommunity-setparameter} additive
```

| Syntax Description | |
|-----------------------------------|-----------------------------------------------------------------------------------------------------------------------------|
| <i>rt-extcommunity-set-name</i> | Route target extended community set name. |
| <i>rt-inline-extcommunity-set</i> | Inline route target extended community set. The inline route target extended community set must be enclosed in parentheses. |
| <i>parameter</i> | 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 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.

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/RP0/CPU0:router(config-rpl)# set extcommunity rt (10:24)
RP/0/RP0/CPU0:router(config-rpl)# set extcommunity rt (10:24, $as:24, $as:$tag)
RP/0/RP0/CPU0:router(config-rpl)# set extcommunity rt (10:24, internet) additive
RP/0/RP0/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 | <p><i>number</i> Value of the precedence. The precedence value can be a number from 0 to 7:</p> <ul style="list-style-type: none"> • 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) |
| | <p><i>parameter</i> Parameter name. The parameter name must be preceded with a “\$.”</p> |

| | |
|------------------------|-------------------------------|
| Command Default | No default behavior or values |
|------------------------|-------------------------------|

| | |
|----------------------|----------------------------|
| Command Modes | Route-policy configuration |
|----------------------|----------------------------|

| | | |
|------------------------|----------------|------------------------------|
| Command History | Release | Modification |
| | Release 5.0.0 | This command was introduced. |

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

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 NCS 6000 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/RP0/CPU0:router(config)# route-policy policy_1
RP/0/RP0/CPU0:router(config-rpl)# set ip-precedence 3
```

```
RP/0/RP0/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 | <i>number</i> 24-bit integer number. Range is from 0 to 16777215. | |
| | <i>parameter</i> 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 5.0.0 | This command was introduced. |
| Usage Guidelines | <p>To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.</p> <p>Use the set isis-metric command to set the IS-IS metric attribute value for routes that are redistributed into IS-IS.</p> | |
| 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/RP0/CPU0:router(config)# route-policy policy_1
RP/0/RP0/CPU0:router(config-rpl)# set isis-metric 1000
RP/0/RP0/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

Release 5.0.0

Modification

This command was introduced.

Usage Guidelines

To use this command, you must be in a user group associated with a task group that 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/RP0/CPU0:router(config)# route-policy labelpolicy
RP/0/RP0/CPU0:router(config-rpl)# if destination in (206.141.1.0/24) then
RP/0/RP0/CPU0:router(config-rpl)# set label explicit-null
RP/0/RP0/CPU0:router(config-rpl)# elseif destination in (206.141.3.0/24) then
RP/0/RP0/CPU0:router(config-rpl)# drop
RP/0/RP0/CPU0:router(config-rpl)# elseif destination in (206.141.4.0/24) then
RP/0/RP0/CPU0:router(config-rpl)# set label explicit-null
RP/0/RP0/CPU0:router(config-rpl)# endif
RP/0/RP0/CPU0:router(config-rpl)# end-policy
```

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 neighbor 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/RP0/CPU0:router#config
RP/0/RP0/CPU0:router(config)#route-policy mordac
RP/0/RP0/CPU0:router(config-rpl)#set large-community (1:2:3, peeras:2:3)
RP/0/RP0/CPU0:router(config-rpl)#end-set
RP/0/RP0/CPU0:router(config)#large-community-set catbert
RP/0/RP0/CPU0:router(config-largecomm)#1:2:3,
RP/0/RP0/CPU0:router(config-largecomm)#5:2:3
RP/0/RP0/CPU0:router(config-largecomm)#end-set
RP/0/RP0/CPU0:router(config)#route-policy wally
RP/0/RP0/CPU0:router(config-rpl)#set large-community catbert additive
RP/0/RP0/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-1-2parameter}
```

| | | |
|---------------------------|------------------|-----------------------------------------------------------------------------------------------|
| 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. |
| | <i>parameter</i> | 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 5.0.0 | 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 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/RP0/CPU0:router(config)# route-policy bgp_isis_redist
RP/0/RP0/CPU0:router(config-rpl)# if destination in (172.2.0.0/16 ge 16) then
RP/0/RP0/CPU0:router(config-rpl)# set level level-2
RP/0/RP0/CPU0:router(config-rpl)# endif
RP/0/RP0/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 | |
|--------------------|------------------------------------------------------------------------|
| <i>number</i> | Value assigned to a 32-bit unsigned integer. Range is 0 to 4294967295. |
| <i>parameter</i> | 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 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. |
|------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

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 NCS 6000 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/RP0/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 | | |
|----------------------|------------------------------------------------------------------------------------------------------------------------|--|
| <i>number</i> | Value assigned to a 32-bit unsigned integer. Range is 0 to 4294967295. | |
| <i>parameter</i> | 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 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.

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/RP0/CPU0:router(config-rpl)# set med 156
RP/0/RP0/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/RP0/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 | rib-metric-as-internal | rib-metric-as-external} parameter
```

| Syntax Description | Parameter | Description |
|--------------------|-------------------------------|------------------------------------------------------------------|
| | internal | Sets metric type to internal. |
| | external | Sets the metric type to external. |
| | rib-metric-as-internal | Uses RIB metric and sets IS-IS internal metric type. |
| | rib-metric-as-external | Uses RIB metric and sets IS-IS external metric type. |
| | <i>parameter</i> | 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 5.0.0 | 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 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/RP0/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. |
| | <i>parameter</i> | 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 5.0.0 | 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 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/RP0/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.

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

Release 5.0.0

Modification

This command was introduced.

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 **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/RP0/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/RP0/CPU0:router(config-rpl)# set next-hop $nexthop
```

set next-hop self in VPN address family

Route reflector does not normally allocate a local label for VPN routes, because it is not next-hop and merely reflects the label from the real next-hop. When the route reflector becomes next-hop, it needs a local label.

With this feature, we will have a route reflector set next-hop-self in a neighbor outbound policy in a VPN address family, either vpnv4 or vpnv6.

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 | egpparameter}
```

| 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). |
| <i>parameter</i> | 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 5.0.0 | 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 **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/RP0/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 | <i>number</i> Value assigned to a 24-bit unsigned integer. Range is 0 to 4294967295. |
| | <i>parameter</i> 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 5.0.0 | 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 **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/RP0/CPU0:router(config)# route-policy policy_1
RP/0/RP0/CPU0:router(config-rpl)# set ospf-metric 1000
RP/0/RP0/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. |
| <i>number</i> | 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 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 NCS 6000 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 install 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
```


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 | <i>number</i> QoS group ID. Range is from 0 to 31. | | | | |
|---------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------|--------------|---------------|------------------------------|
| | <i>parameter</i> Parameter name. The parameter name must be preceded with a "\$." | | | | |
| Command Default | No default behavior or values | | | | |
| Command Modes | Route-policy configuration | | | | |
| Command History | <table border="1"> <thead> <tr> <th>Release</th> <th>Modification</th> </tr> </thead> <tbody> <tr> <td>Release 5.0.0</td> <td>This command was introduced.</td> </tr> </tbody> </table> | Release | Modification | Release 5.0.0 | This command was introduced. |
| Release | Modification | | | | |
| 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.

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 NCS 6000 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/RP0/CPU0:router(config)# route-policy policy_1
RP/0/RP0/CPU0:router(config-rpl)# set qos-group 12
RP/0/RP0/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 | <i>number</i> Value assigned to a 32-bit unsigned integer. Range is 0 to 4294967295. |
| | <i>parameter</i> 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 5.0.0 | 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 **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/RP0/CPU0:router(config)# route-policy policy_1
RP/0/RP0/CPU0:router(config-rpl)# set rib-metric 1000
RP/0/RP0/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 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.

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/RP0/CPU0:router(config)# route-policy policy_1
RP/0/RP0/CPU0:router(config-rpl)# set rip-metric 10
RP/0/RP0/CPU0:router(config-rpl)# end-policy
```

Related Commands

| Command | Description |
|------------------------------------|--------------------------------------|
| add , on page 1179 | 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 | <i>number</i> Value assigned to a 16-bit unsigned integer. Range is from 0 to 65535. |
| | <i>parameter</i> 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 5.0.0 |

| | |
|-------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Usage Guidelines | To use this command, you must be in a user group associated with a task group that 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/RP0/CPU0:router(config)# route-policy policy_1
RP/0/RP0/CPU0:router(config-rpl)# set rip-tag 1000
RP/0/RP0/CPU0:router(config-rpl)# end-policy
```

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 XR Config 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 | XR Config | | | | |
| Usage Guidelines | To use this command, you must be in a user group associated with a task 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 | <table border="1"> <thead> <tr> <th>Task ID</th> <th>Operations</th> </tr> </thead> <tbody> <tr> <td>route-policy</td> <td>read</td> </tr> </tbody> </table> | Task ID | Operations | route-policy | read |
| Task ID | Operations | | | | |
| route-policy | read | | | | |

Examples

Based on service agreement, if a user AVPN advertises VPN prefixes with RT-set {1:1, 2:2} and user Team10 advertises VPN prefixes with RTs {3:3, 4:4}. On the DUT, the user wants to limit AVPN prefixes to 100, Team10 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 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-policy | read, write |

This example sets SPF priority as critical:

```
RP/0/RP0/CPU0:router#configure
RP/0/RP0/CPU0:router(config)#route-policy policy_spf_priority
RP/0/RP0/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 | <i>number</i> Value assigned to a 32-bit unsigned integer. Range is from 0 to 4294967295. | |
| | <i>parameter</i> 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 5.0.0 | 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 **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/RP0/CPU0:router(config-rpl)# set tag 10
```

In this example, the tag attribute is set to a parameter value \$tag_param:

```
RP/0/RP0/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

Release 5.0.0

Modification

This command was introduced.

Usage Guidelines

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

Release 5.0.0

Modification

This command was introduced.

Usage Guidelines

To use this command, you must be in a user group associated with a task group that 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/RP0/CPU0:router(config-rpl)# set weight 10  
RP/0/RP0/CPU0:router(config-rpl)# set weight $weight_param
```

show rpl

To display system-wide RPL configuration, use the **show rpl** command in XR 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 <i>configuration-limit</i> | (Optional) | Displays the number of lines to which configuration is limited. Range is 1 to 131072. The <i>configuration-limit</i> argument is available if the running-config keyword is specified. |
| policies <i>policies-limit</i> | (Optional) | Displays the limit on the number of policies. Range is 1 to 5000. The <i>configuration-limit</i> 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 XR EXEC

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | 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 | Operations |
|---------|--------------|----------------|
| | route-policy | read, write |

Examples

The following shows the output of the `show running-config rpl` command:

```
RP/0/RP0/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 1388 | 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 XR 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 | XR EXEC |
|----------------------|---------|

| Command History | Release | Modification |
|------------------------|----------------|------------------------------|
| | 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. |
|-------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

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/RP0/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/RP0/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/RP0/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 1341 | 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 1344 | 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 1347 | 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 1347 | 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 XR 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 | XR EXEC |
|----------------------|---------|

| Command History | Release | Modification |
|------------------------|----------------|------------------------------|
| | 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. |
|-------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

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/RP0/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,
```

show rpl active community-set

```

    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/RP0/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/RP0/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 1338 | 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 1344 | 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 1347 | 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 1347 | 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 1350 | 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 XR 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

XR EXEC

Command History

| Release | Modification |
|---------------|------------------------------|
| 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.

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/RP0/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/RP0/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:

show rpl active extcommunity-set

```
RP/0/RP0/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 1338 | 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 1341 | 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 1347 | 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 1347 | 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 1350 | 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 XR 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 | XR EXEC | |
| Command History | Release | Modification |
| | Release 5.0.0 | This command was introduced. |
| Usage Guidelines | <p>To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.</p> <p>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.</p> | |
| 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/RP0/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,
```

show rpl active prefix-set

```

    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/RP0/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/RP0/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 1338 | 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 1341 | 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 1344 | 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 1406 | 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 1350 | 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 XR 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 | XR EXEC |
|----------------------|---------|

| Command History | Release | Modification |
|------------------------|----------------|------------------------------|
| | 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. |
|-------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

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/RP0/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 1338 | 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 1341 | 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 1344 | 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 1347 | 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 1347 | 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 XR 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 | XR EXEC |
|----------------------|---------|

| Command History | Release | Modification |
|------------------------|----------------|------------------------------|
| | 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. |
|-------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

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/RP0/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/RP0/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/RP0/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 1338 | 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 1341 | 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 1344 | 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 1347 | 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 1350 | 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 XR 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

XR EXEC

Command History

Release

Release 5.0.0

Modification

This command was introduced.

Usage Guidelines

To use this command, you must be in a user group associated with a task group that 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/RP0/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/RP0/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 1361 | Displays the configuration of a named community set. |
| show rpl extcommunity-set, on page 1368 | Displays the configuration of a named extended community set. |
| show rpl route-policy, on page 1403 | Displays the configuration of a named route policy. |
| show rpl prefix-set, on page 1392 | 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 XR EXEC mode.

```
show rpl as-path-set name attachpoints
```

| | |
|---------------------------|-------------------------------------|
| Syntax Description | <i>name</i> Name of an AS path set. |
|---------------------------|-------------------------------------|

| | |
|------------------------|-------------------------------|
| Command Default | No default behavior or values |
|------------------------|-------------------------------|

| | |
|----------------------|---------|
| Command Modes | XR EXEC |
|----------------------|---------|

| Command History | Release | Modification |
|------------------------|----------------|------------------------------|
| | 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. |
|-------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

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/RP0/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/RP0/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/RP0/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 145: 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 1363 | Displays all the policies used at an attach point that reference the named community set. |
| show rpl route-policy attachpoints, on page 1406 | Displays all the policies used at an attach point that reference the named policy. |
| show rpl prefix-set attachpoints, on page 1394 | 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 XR EXEC mode.

```
show rpl as-path-set name references [brief]
```

| | |
|---------------------------|-------------------------------------------------------------------------------------------------------------------------------|
| Syntax Description | <i>name</i> 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 | XR EXEC |
|----------------------|---------|

| | | |
|------------------------|----------------|------------------------------|
| Command History | Release | Modification |
| | 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.

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/RP0/CPU0:router# show rpl route-policy policy_1

route-policy policy_1
 if (destination in prefix_set_ex1) then
  set local-preference 100
 endif
```

show rpl as-path-set references

```

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/RP0/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 146: 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 1365 | Lists all policies that reference the named community set. |
| show rpl route-policy references, on page 1410 | Lists all policies that reference the named policy. |
| show rpl prefix-set references, on page 1397 | 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 XR 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

XR EXEC

Command History

Release

Release 5.0.0

Modification

This command was introduced.

Usage Guidelines

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

show rpl community-set

```

    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/RP0/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 1354 | Displays the configuration of a named AS path set. |
| show rpl extcommunity-set, on page 1368 | Displays the configuration of a named extended community set. |
| show rpl prefix-set, on page 1392 | Displays the configuration of a named prefix set. |
| show rpl rd-set, on page 1399 | Displays the configuration of a named RD set. |
| show rpl route-policy, on page 1403 | 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 XR EXEC mode.

```
show rpl community-set name attachpoints
```

| | |
|---------------------------|--------------------------------------|
| Syntax Description | <i>name</i> Name of a community set. |
|---------------------------|--------------------------------------|

| | |
|------------------------|-------------------------------|
| Command Default | No default behavior or values |
|------------------------|-------------------------------|

| | |
|----------------------|---------|
| Command Modes | XR EXEC |
|----------------------|---------|

| Command History | Release | Modification |
|------------------------|----------------|------------------------------|
| | 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.

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/RP0/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 147: 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 1356 | Displays all the policies used at an attach point that reference the named AS path set. |
| show rpl prefix-set attachpoints, on page 1394 | Displays all the policies used at an attach point that reference the named prefix set. |
| show rpl route-policy attachpoints, on page 1406 | 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 XR EXEC mode.

show rpl community-set *name* **references** [**brief**]

| Syntax Description | <p><i>name</i> Name of a community set.</p> <hr/> <p>brief (Optional) Limits the output to just the summary table and not the detailed information for the community set.</p> | | | | |
|---------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------|--------------|---------------|------------------------------|
| Command Default | No default behavior or values | | | | |
| Command Modes | XR EXEC | | | | |
| Command History | <table border="1"> <thead> <tr> <th>Release</th> <th>Modification</th> </tr> </thead> <tbody> <tr> <td>Release 5.0.0</td> <td>This command was introduced.</td> </tr> </tbody> </table> | Release | Modification | Release 5.0.0 | This command was introduced. |
| Release | Modification | | | | |
| Release 5.0.0 | This command was introduced. | | | | |
| Usage Guidelines | <p>To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.</p> <p>Use the show rpl community-set references command to display all the policies that reference the named community set.</p> <p>Use the optional brief keyword to limit the output to just a summary table and not the detailed information for the community set.</p> | | | | |
| Task ID | <table border="1"> <thead> <tr> <th>Task ID</th> <th>Operations</th> </tr> </thead> <tbody> <tr> <td>route-policy</td> <td>read</td> </tr> </tbody> </table> | Task ID | Operations | route-policy | read |
| 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
```

show rpl community-set references

```

    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/RP0/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 148: 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 1359 | Lists all policies that reference the named AS path set. |
| show rpl prefix-set references, on page 1397 | Lists all policies that reference the named prefix set. |

| Command | Description |
|----------------------------------------------------------------|-----------------------------------------------------|
| show rpl rd-set references, on page 1401 | Lists all policies that reference the named RD set. |
| show rpl route-policy references, on page 1410 | 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 XR EXEC mode.

```
show rpl extcommunity-set [name [{attachpoints | references}]] [{cost | rt | soo}] [name] [brief]
[states]
```

Syntax Description

| | |
|---------------------|-------------------------------------------------------------------------------------------------------------------|
| <i>name</i> | (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

XR EXEC

Command History

| Release | Modification |
|---------------|------------------------------|
| 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.

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/RP0/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/RP0/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 extsool
  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/RP0/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,
```

show rpl extcommunity-set

```

    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 1354 | Displays the configuration of a named AS path set. |
| show rpl community-set, on page 1361 | Displays the configuration of a named community set. |
| show rpl prefix-set, on page 1392 | Displays the configuration of a named prefix set. |
| show rpl rd-set, on page 1399 | Displays the configuration of a named RD set. |
| show rpl route-policy, on page 1403 | 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 XR 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 | XR EXEC |
|----------------------|---------|

| Command History | Release | Modification |
|------------------------|----------------|------------------------------|
| | 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.

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
!
```

show rpl inactive as-path-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
!
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/RP0/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 1374 | 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 1377 | 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 1380 | 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 1383 | 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 1385 | 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 XR 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 | XR EXEC |
|----------------------|---------|

| Command History | Release | Modification |
|------------------------|----------------|------------------------------|
| | 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. |
|-------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

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/RP0/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 1371 | 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 1377 | 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 1380 | 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 1383 | 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 1385 | 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 XR 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 | XR EXEC |
|----------------------|---------|

| Command History | Release | Modification |
|------------------------|----------------|------------------------------|
| | 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.

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/RP0/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 1371 | 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 1374 | 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 1380 | 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 1383 | 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 1385 | 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 XR 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 | XR EXEC |
|----------------------|---------|

| Command History | Release | Modification |
|------------------------|----------------|------------------------------|
| | 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. |
|-------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

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/RP0/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 1371 | 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 1374 | 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 1377 | 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 1385 | 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 1383 | 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 XR 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 | XR EXEC | |
| Command History | Release | Modification |
| | Release 5.0.0 | This command was introduced. |
| Usage Guidelines | <p>To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.</p> <p>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.</p> | |
| 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/RP0/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 1371 | 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 1374 | 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 1377 | 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 1380 | 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 1385 | 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 XR 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 | XR EXEC |
|----------------------|---------|

| Command History | Release | Modification |
|------------------------|----------------|------------------------------|
| | 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.

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/RP0/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 1371 | 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 1374 | 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 1377 | 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 1380 | 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 1383 | 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 XR 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 XR EXEC

| Command History | Release | Modification |
|------------------------|----------------|------------------------------|
| | 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.

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/RP0/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 149: show rpl maximum Field Descriptions, on page 1389](#) describes the significant fields shown in the display.

Table 149: 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 1294 | 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 XR 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 | XR EXEC |
|----------------------|---------|

| Command History | Release | Modification |
|------------------------|----------------|------------------------------|
| | 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 | 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/RP0/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 XR 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

XR EXEC

Command History

Release

Release 5.0.0

Modification

This command was introduced.

Usage Guidelines

To use this command, you must be in a user group associated with a task 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/RP0/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 1354 | Displays the configuration of a named AS path set. |

| Command | Description |
|---------------------------------------------------------|---------------------------------------------------------------|
| show rpl community-set, on page 1361 | Displays the configuration of a named community set. |
| show rpl extcommunity-set, on page 1368 | Displays the configuration of a named extended community set. |
| show rpl route-policy, on page 1403 | 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 XR EXEC mode.

show rpl prefix-set *name* attachpoints

| | |
|---------------------------|-----------------------------------|
| Syntax Description | <i>name</i> Name of a prefix set. |
|---------------------------|-----------------------------------|

| | |
|------------------------|-------------------------------|
| Command Default | No default behavior or values |
|------------------------|-------------------------------|

| | |
|----------------------|---------|
| Command Modes | XR EXEC |
|----------------------|---------|

| Command History | Release | Modification |
|------------------------|----------------|------------------------------|
| | 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. |
|-------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

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/RP0/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 150: 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 1356 | Displays all the policies used at an attach point that reference the named AS path set. |
| show rpl community-set attachpoints, on page 1363 | Displays all the policies used at an attach point that reference the named community set. |

| Command | Description |
|------------------------------------------------------------------|------------------------------------------------------------------------------------|
| show rpl route-policy attachpoints, on page 1406 | 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 XR EXEC mode.

```
show rpl prefix-set name references [brief]
```

| | |
|---------------------------|------------------------------------------------------------------------------------------------------------------------------|
| Syntax Description | <i>name</i> 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 | XR EXEC |
|----------------------|---------|

| | | |
|------------------------|----------------|------------------------------|
| Command History | Release | Modification |
| | 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.

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/RP0/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 151: 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 1359 | Lists all the policies that reference the named AS path set. |
| show rpl community-set references, on page 1365 | Lists all the policies that reference the named community set. |
| show rpl route-policy references, on page 1410 | 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 XR 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 XR EXEC

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | 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.

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/RP0/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 1354 | Displays the configuration of a named AS path set. |
| | show rpl community-set, on page 1361 | Displays the configuration of a named community set. |

| Command | Description |
|---------------------------------------------------------|---------------------------------------------------------------|
| show rpl extcommunity-set, on page 1368 | Displays the configuration of a named extended community set. |
| show rpl prefix-set, on page 1392 | Displays the configuration of a named prefix set. |
| show rpl route-policy, on page 1403 | Displays the configuration of a named route 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 XR EXEC mode.

```
show rpl rd-set name references [brief]
```

| Syntax Description | <i>name</i> 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 | XR EXEC | | | | |
| Command History | <table border="1"> <thead> <tr> <th>Release</th> <th>Modification</th> </tr> </thead> <tbody> <tr> <td>Release 5.0.0</td> <td>This command was introduced.</td> </tr> </tbody> </table> | Release | Modification | Release 5.0.0 | This command was introduced. |
| Release | Modification | | | | |
| 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.

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/RP0/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 152: 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 1359 | Lists all the policies that reference the named AS path set. |
| show rpl community-set references, on page 1365 | Lists all the policies that reference the named community set. |
| show rpl prefix-set references, on page 1397 | Lists all the policies that reference the named prefix set. |
| show rpl route-policy references, on page 1410 | 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 XR 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 | XR EXEC |
|----------------------|---------|

| | | |
|------------------------|----------------|------------------------------|
| Command History | Release | Modification |
| | 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.

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/RP0/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
```

show rpl route-policy

```

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/RP0/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 1354 | Displays the configuration of a named AS path set. |
| show rpl community-set, on page 1361 | Displays the configuration of a named community set. |
| show rpl extcommunity-set, on page 1368 | Displays the configuration of a named extended community set. |
| show rpl prefix-set, on page 1392 | 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 XR EXEC mode.

show rpl route-policy *name* attachpoints

| | |
|---------------------------|-------------------------------|
| Syntax Description | <i>name</i> Name of a policy. |
|---------------------------|-------------------------------|

| | |
|------------------------|-------------------------------|
| Command Default | No default behavior or values |
|------------------------|-------------------------------|

| | |
|----------------------|---------|
| Command Modes | XR EXEC |
|----------------------|---------|

| | | |
|------------------------|----------------|------------------------------|
| Command History | Release | Modification |
| | 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. |
|-------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

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/RP0/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/RP0/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
!

```

| Related Commands | Command | Description |
|------------------|-------------------------------------------------------------------|-------------------------------------------------------------------------------------------|
| | show rpl as-path-set attachpoints, on page 1356 | Displays all the policies used at an attach point that reference the named AS path set. |
| | show rpl community-set attachpoints, on page 1363 | Displays all the policies used at an attach point that reference the named community set. |
| | show rpl prefix-set attachpoints, on page 1394 | 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 XR EXEC mode.

show rpl route-policy *name* **inline**

| | |
|---------------------------|-------------------------------|
| Syntax Description | <i>name</i> Name of a policy. |
|---------------------------|-------------------------------|

| | |
|------------------------|-------------------------------|
| Command Default | No default behavior or values |
|------------------------|-------------------------------|

| | |
|----------------------|---------|
| Command Modes | XR EXEC |
|----------------------|---------|

| Command History | Release | Modification |
|------------------------|----------------|------------------------------|
| | 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. |
|-------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

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/RP0/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/RP0/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 XR EXEC mode.

show rpl route-policy *name* **references** [**brief**]

| | |
|---------------------------|--------------------------------------------------------------------------------------------------------------------------|
| Syntax Description | <i>name</i> 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 XR EXEC

| Command History | Release | Modification |
|------------------------|----------------|------------------------------|
| | 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.

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/RP0/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 153: 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 1359 | Lists all policies that reference the named AS path set. |
| show rpl community-set references, on page 1365 | Lists all policies that reference the named community set. |
| show rpl prefix-set references, on page 1397 | 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 XR EXEC mode.

show rpl route-policy *name* **uses** {**policies** | **sets** | **all**} [**direct**]

| Syntax Description | |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>name</i> | 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. |
| <i>direct</i> | (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 XR EXEC

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | 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.

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/RP0/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/RP0/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 154: 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 XR 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 | XR EXEC |
|----------------------|---------|

| Command History | Release | Modification |
|------------------------|----------------|------------------------------|
| | 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. |
|-------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

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:

show rpl unused as-path-set

```
RP/0/RP0/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 1419 | Displays the community sets that are not referenced at all. |
| show rpl unused extcommunity-set, on page 1422 | Displays the extended community sets that are not referenced at all. |
| show rpl unused prefix-set, on page 1424 | Displays the prefix sets that are not referenced at all. |
| show rpl unused rd-set, on page 1427 | Displays the RD sets that are not referenced at all. |
| show rpl unused route-policy, on page 1429 | 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 XR 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 | XR EXEC | |
| Command History | Release | Modification |
| | Release 5.0.0 | This command was introduced. |
| Usage Guidelines | <p>To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.</p> <p>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.</p> | |
| Task ID | Task ID | Operations |
| | route-policy | read |
| Examples | <p>This example shows the following sample configuration:</p> <pre> 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 </pre> | |

```

!
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/RP0/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 1416 | Displays the AS path sets that are not referenced at all. |
| show rpl unused extcommunity-set, on page 1422 | Displays the extended community sets that are not referenced at all. |
| show rpl unused prefix-set, on page 1424 | Displays the prefix sets that are not referenced at all. |
| show rpl unused rd-set, on page 1427 | Displays the RD sets that are not referenced at all. |
| show rpl unused route-policy, on page 1429 | 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 XR 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

XR EXEC

Command History

| Release | Modification |
|---------------|------------------------------|
| 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.

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/RP0/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 1416 | Displays the AS path sets that are not referenced at all. |
| show rpl unused community-set, on page 1419 | Displays community sets that are not referenced at all. |
| show rpl unused prefix-set, on page 1424 | Displays prefix sets that are not referenced at all. |
| show rpl unused rd-set, on page 1427 | Displays the RD sets that are not referenced at all. |
| show rpl unused route-policy, on page 1429 | 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 XR 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 | XR EXEC |
|----------------------|---------|

| Command History | Release | Modification |
|------------------------|----------------|------------------------------|
| | 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. |
|-------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

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

```

show rpl unused prefix-set

```

        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/RP0/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 1416 | Displays AS path sets that are not referenced at all. |
| show rpl unused community-set, on page 1419 | Displays community sets that are not referenced at all. |
| show rpl unused extcommunity-set, on page 1422 | Displays extended community sets that are not referenced at all. |
| show rpl unused rd-set, on page 1427 | Displays the RD sets that are not referenced at all. |
| show rpl unused route-policy, on page 1429 | 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 XR 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 | XR EXEC |
|----------------------|---------|

| Command History | Release | Modification |
|------------------------|----------------|------------------------------|
| | 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. |
|-------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

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 |

| | |
|-----------------|-------------------------------------------------------------------------------|
| Examples | The show rpl unused rd-set command displays the following information: |
|-----------------|-------------------------------------------------------------------------------|

```
RP/0/RP0/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 1416 | Displays the AS path sets that are not referenced at all. |
| | show rpl unused community-set, on page 1419 | Displays the community sets that are not referenced at all. |

| Command | Description |
|----------------------------------------------------------------|----------------------------------------------------------------------|
| show rpl unused extcommunity-set, on page 1422 | Displays the extended community sets that are not referenced at all. |
| show rpl unused prefix-set, on page 1424 | Displays the prefix sets that are not referenced at all. |
| show rpl unused route-policy, on page 1429 | 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 XR 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 | XR EXEC | |
| Command History | Release | Modification |
| | Release 5.0.0 | This command was introduced. |
| Usage Guidelines | <p>To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.</p> <p>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.</p> | |
| Task ID | Task ID | Operations |
| | route-policy | read |

Examples

This example shows the following sample configuration:

```
RP/0/RP0/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/RP0/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 1416 | Displays AS path sets that are not referenced at all. |
| show rpl unused community-set, on page 1419 | Displays community sets that are not referenced at all. |
| show rpl unused extcommunity-set, on page 1422 | Displays extended community sets that are not referenced at all. |
| show rpl unused prefix-set, on page 1424 | Displays prefix sets that are not referenced at all. |
| show rpl unused rd-set, on page 1427 | 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-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 name. The parameter name must be preceded with a "\$."

Command Default

No default behavior or values

Command Modes

Route-policy configuration

Command History

Release

Release 5.0.0

Modification

This command was introduced.

Usage Guidelines

To use this command, you must be in a user group associated with a task group that 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/RP0/CPU0:router(config)# route-policy policy-A
RP/0/RP0/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/RP0/CPU0:router(config)# route-policy policy-B
RP/0/RP0/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/RP0/CPU0:router(config)# route-policy policy-C
RP/0/RP0/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 1277 | 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 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.

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 1440](#) 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/RP0/CPU0:router(config)# oute-policy check-aggregator
RP/0/RP0/CPU0:router(config-rpl)# if destination in (10.1.0.0/16) then
RP/0/RP0/CPU0:router(config-rpl-if)# suppress-route

RP/0/RP0/CPU0:router(config-rpl-if)# endif
RP/0/RP0/CPU0:router(config-rpl-if)# end-policy
```

| Related Commands | Command | Description |
|------------------|------------------------------------------------|-------------------------------------------------------------------------|
| | unsuppress-route, on page 1440 | 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. |
| <i>integer</i> | Integer value. Range is 0 to 4294967295. |
| <i>parameter</i> | 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 5.0.0 | 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 **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-name|inline-tag-set}parameter}
```

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 name. The parameter name must be preceded with a "\$."

parameter

Command Default

No default behavior or values

Command Modes

Route-policy configuration

Command History

Release

Release 5.0.0

Modification

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/RP0/CPU0:router (config) #tag-set my-tag-set
RP/0/RP0/CPU0:router (config-tag) #1000
```

```
RP/0/RP0/CPU0:router(config-tag)#3000
RP/0/RP0/CPU0:router(config-tag)#end-set

RP/0/RP0/CPU0:router(config)#route-policy use-tag-in
RP/0/RP0/CPU0:router(config-rpl)#if tag in my-tag-set then
RP/0/RP0/CPU0:router(config-rpl-if)#set local-preference 100
RP/0/RP0/CPU0:router(config-rpl-if)#elseif tag in (2000, 4000) then
RP/0/RP0/CPU0:router(config-rpl-elseif)#set local-preference 200
RP/0/RP0/CPU0:router(config-rpl-elseif)#endif
RP/0/RP0/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 XR Config mode. To remove a named tag set, use the **no** form of this command.

```
tag-set name
no tag-set name
```

| | | |
|---------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------|
| Syntax Description | <i>name</i> | Name of a tag set. |
| Command Default | None | |
| Command Modes | XR Config | |
| Command History | Release | Modification |
| | Release 5.0.0 | 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/RP0/CPU0:router(config)#tag-set my-tag-set
RP/0/RP0/CPU0:router(config-tag)#1000
RP/0/RP0/CPU0:router(config-tag)#3000
RP/0/RP0/CPU0:router(config-tag)#end-set

RP/0/RP0/CPU0:router(config)#route-policy use-tag-in
RP/0/RP0/CPU0:router(config-rpl)#if tag in my-tag-set then
RP/0/RP0/CPU0:router(config-rpl-if)#set local-preference 100
RP/0/RP0/CPU0:router(config-rpl-if)#elseif tag in (2000, 4000) then
RP/0/RP0/CPU0:router(config-rpl-elseif)#set local-preference 200
RP/0/RP0/CPU0:router(config-rpl-elseif)#endif
RP/0/RP0/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 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.

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/RP0/CPU0:router(config)# route-policy check-aggregate
RP/0/RP0/CPU0:router(config-rpl)# if destination in (10.1.0.0/16) then
RP/0/RP0/CPU0:router(config-rpl-if)# unsuppress-route

RP/0/RP0/CPU0:router(config-rpl-if)# endif
RP/0/RP0/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 1435 | Indicates that a given component of a BGP aggregate should be suppressed. |

■ `unsuppress-route`



Static Routing Commands

This module describes the commands used to establish static routes.

For detailed information about static routing concepts, configuration tasks, and examples, see the *Implementing Static Routes on Software* module in the *Routing Configuration Guide for Cisco NCS 6000 Series Routers*.

- [address-family \(static\)](#), on page 1444
- [maximum path \(static\)](#), on page 1445
- [route \(static\)](#), on page 1447
- [router static](#), on page 1450

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}
no address-family {ipv4 | ipv6} {unicast}
```

| 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. |

| Command Modes | Router static configuration |
|---------------|-----------------------------|
|---------------|-----------------------------|

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | 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.

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 | Operations |
|---------|--------|----------------|
| | static | read, write |

Examples

The following example shows how to enter IPv6 unicast address family mode:

```
RP/0/RP0/CPU0:router(config)# router static
RP/0/RP0/CPU0:router(config-static)# address-family ipv6 unicast
RP/0/RP0/CPU0:router(config-static-afi)#
```

Related Commands

| Command | Description |
|----------------------------------------------|-----------------------------|
| route (static), on page 1447 | 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. |
| | <i>value</i> | Maximum number of static routes for the given AFI. The range is 1 to 140000. |

| | |
|------------------------|---------------------|
| Command Default | <i>value</i> : 4000 |
|------------------------|---------------------|

| | |
|----------------------|-----------------------------|
| Command Modes | Static router configuration |
|----------------------|-----------------------------|

| | | |
|------------------------|----------------|------------------------------|
| Command History | Release | Modification |
| | 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 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/RP0/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/RP0/CPU0:router(config-static)# no maximum path ipv4 100000
```

Related Commands

| Command | Description |
|----------------------------------------------|------------------------------------------------|
| route (static), on page 1447 | Enters static router configuration mode. |
| show route static | Displays the static routes in a routing table. |

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 {ip-address | type interface-path-id [{ip-address | type interface-path-id}] [track
track-object-name] [ ] [distance] [description text] [tag tag] [permanent]}
no prefix/mask {ip-address | type interface-path-id [{ip-address | type interface-path-id}] [track
track-object-name] [ ] [distance] [description text] [tag tag] [permanent]}
```

Syntax Description

| | |
|--------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>prefix / mask</i> | IP route prefix and prefix mask for the destination. The network mask can be specified in either of two ways: <ul style="list-style-type: none"> • 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. |
| <i>ip-address</i> | IP address of the next hop that can be used to reach that network. <ul style="list-style-type: none"> • 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. <p>Note A forwarding router's IP address or an interface or virtual interface path ID can be configured, in any order.</p> |
| <i>type</i> | Interface type. For more information, use the question mark (?) online help function. |
| <i>interface-path-id</i> | Physical interface or virtual interface. <p>Note Use the show interfaces command to see a list of all interfaces currently configured on the router.</p> <p>For more information about the syntax for the router, use the question mark (?) online help function.</p> <p>Note A forwarding router's IP address or an interface or virtual interface path ID can be configured, in any order.</p> |
| <i>distance</i> | (Optional) Administrative distance. Range is 1 to 254. |
| description <i>text</i> | (Optional) Specifies a description of the static route. |

| | |
|------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------|
| tag <i>tag</i> | (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 <i>track-object-name</i> | Enables object tracking for static route. |

Command Default No static route is established.

Command Modes Address family configuration

| Command History | Release | Modification |
|------------------------|----------------|------------------------------|
| | 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.

A static route is appropriate when the software cannot dynamically build a route to the destination.

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.

| Task ID | Task ID | Operations |
|----------------|----------------|-------------------|
| | static | read, write |

Examples

The following example shows how to configure IPv6 unicast address family static routes:

```
RP/0/RP0/CPU0:router(config)# router static
RP/0/RP0/CPU0:router(config-static)# address-family ipv6 unicast
RP/0/RP0/CPU0:router(config-static-afi)# 2b11::327a:7b00/120 GigabitEthernet0/2/0/7
RP/0/RP0/CPU0:router(config-static-afi)# 2b11::327a:7b00/120 GigabitEthernet0/6/0/0
RP/0/RP0/CPU0:router(config-static-afi)# 2b11::327a:7b00/120 2b11::2f01:4c
RP/0/RP0/CPU0:router(config-static-afi)# 2b11::327a:7b00/120 2b11::2f01:4d
RP/0/RP0/CPU0:router(config-static-afi)# 2b11::327a:7b00/120 2b11::2f01:4e
RP/0/RP0/CPU0:router(config-static-afi)# 2b11::327a:7b00/120 2b11::2f01:4f
RP/0/RP0/CPU0:router(config-static-afi)# 2b11::327a:7b00/120 2b11::2f01:50
```

Related Commands

| Command | Description |
|-------------------------------------------------------|-----------------------------------------------------------------------|
| address-family (static), on page 1444 | Enters address family configuration mode. |
| 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. |
| router static, on page 1450 | Enters router static configuration mode. |

router static

To enter static router configuration mode, use the **router static** command in XR Configmode. 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 XR Config

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | 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 | Operations |
|---------|-----------------------------------------|----------------|
| | static | read, write |
| | bgp, ospf, hsrp, isis, vrrp, or network | read, write |

Examples

The following example shows how to enter static router configuration mode:

```
RP/0/RP0/CPU0:router(config)# router static
RP/0/RP0/CPU0:router(config-static)#
```

Related Commands

| Command | Description |
|-------------------------------------------------------|-----------------------------------------------------------------------|
| address-family (static), on page 1444 | 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 1447 | Establishes 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 NCS 6000 Series Routers*.

- [router-convergence](#), on page 1452
- [monitor-convergence \(IS-IS\)](#), on page 1453
- [monitor-convergence \(OSPF\)](#), on page 1454
- [collect-diagnostics \(RCMD\)](#), on page 1455
- [event-buffer-size \(RCMD\)](#), on page 1457
- [max-events-stored \(RCMD\)](#), on page 1458
- [monitoring-interval \(RCMD\)](#), on page 1459
- [node disable \(RCMD\)](#), on page 1460
- [priority \(RCMD\)](#), on page 1462
- [protocol \(RCMD\)](#), on page 1464
- [storage-location](#), on page 1465

router-convergence

To configure route convergence monitoring and enter router convergence monitoring and diagnostics (rcmd) configuration mode, use the **router-convergence** command in XR Config 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 | XR Config |
|----------------------|-----------|

| Command History | Release | Modification |
|------------------------|----------------|------------------------------|
| | 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 |
|----------------|----------------|------------------|
| | rcmd | read, write |

This example shows how to configure router-convergence command and enable rcmd configuration mode:

```
RP/0/RP0/CPU0:router#configure
RP/0/RP0/CPU0:router (config)#router-convergence
RP/0/RP0/CPU0:router (config-rcmd)#
```

| Related Commands | Command | Description |
|-------------------------|-----------------------------------------------------------|----------------------------------------------------------|
| | monitor-convergence (IS-IS), on page 1453 | Enables route convergence monitoring for IS-IS protocol. |
| | monitor-convergence (OSPF), on page 1454 | 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

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | 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 route convergence monitoring for IS-IS under IPv6 unicast SAFI:

```
RP/0/RP0/CPU0:router#configure
RP/0/RP0/CPU0:router(config)#router isis isp
RP/0/RP0/CPU0:router(config-isis)#address-family ipv6 unicast
RP/0/RP0/CPU0:router(config-isis-af)#monitor-convergence
```

| Related Commands | Command | Description |
|------------------|----------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------|
| | router-convergence, on page 1452 | Configures route convergence monitoring and enters router convergence monitoring and diagnostics (rcmd) configuration mode. |
| | monitor-convergence (OSPF), on page 1454 | 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 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 |
|---------|---------|----------------|
| | ospf | read, write |

This example shows how to enable route convergence monitoring for an OSPF process:

```
RP/0/RP0/CPU0:router#configure
RP/0/RP0/CPU0:router (config)#router ospf 100
RP/0/RP0/CPU0:router (config-ospf)#monitor-convergence
```

| Related Commands | Command | Description |
|------------------|-----------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------|
| | router-convergence, on page 1452 | Configures route convergence monitoring and enters router convergence monitoring and diagnostics (rcmd) configuration mode. |
| | monitor-convergence (IS-IS), on page 1453 | 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 | <i>location</i> Specifies the line-card location. |
|---------------------------|---------------------------------------------------|

| | |
|------------------------|-------------------------------------|
| Command Default | Diagnostics collection is disabled. |
|------------------------|-------------------------------------|

| | |
|----------------------|----------------------------------|
| Command Modes | Router-convergence configuration |
|----------------------|----------------------------------|

| | | |
|------------------------|----------------|------------------------------|
| Command History | Release | Modification |
| | 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. |
|-------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

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/RP0/CPU0:router#configure  
RP/0/RP0/CPU0:router#router-convergence  
RP/0/RP0/CPU0:router (config-rcmd) #collect-diagnostics 0/3/CPU0
```

| Related Commands | Command | Description |
|------------------|--------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------|
| | router-convergence, on page 1452 | 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 | <i>number</i> 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 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.

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/RP0/CPU0:router#configure
RP/0/RP0/CPU0:router (config)#router-convergence
RP/0/RP0/CPU0:router (config-rcmd)#event-buffer-size 500
```

| | | |
|-------------------------|--------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------|
| Related Commands | Command | Description |
| | router-convergence, on page 1452 | 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 | <i>number</i> 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 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. |
|-------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

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/RP0/CPU0:router#configure
RP/0/RP0/CPU0:router(config)#router-convergence
RP/0/RP0/CPU0:router(config-rcmd)#max-events-stored 500
```

| Related Commands | Command | Description |
|-------------------------|--------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------|
| | router-convergence, on page 1452 | 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 | <i>minutes</i> 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 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. |
|-------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

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/RP0/CPU0:router#configure
RP/0/RP0/CPU0:router (config)#router-convergence
RP/0/RP0/CPU0:router (config-rcmd)#monitoring-interval 5
```

| Related Commands | Command | Description |
|-------------------------|--------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------|
| | router-convergence, on page 1452 | 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 | <i>node-id</i> 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 | <table border="1"> <thead> <tr> <th>Release</th> <th>Modification</th> </tr> </thead> <tbody> <tr> <td>Release 5.0.0</td> <td>This command was introduced.</td> </tr> </tbody> </table> | Release | Modification | Release 5.0.0 | This command was introduced. |
| Release | Modification | | | | |
| 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.

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/RP0/CPU0:router#configure
RP/0/RP0/CPU0:router(config)#router-convergence
RP/0/RP0/CPU0:router(config-rcmd)#node 0/*/* disable
```

| Related Commands | Command | Description |
|------------------|--------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------|
| | router-convergence, on page 1452 | Configures route convergence monitoring and enters router convergence monitoring and diagnostics (rcmd) configuration mode. |

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. |
| <i>leaf-network-number</i> | 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. |
| <i>value</i> | 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 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.

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/RP0/CPU0:router#configure
RP/0/RP0/CPU0:router(config)#router-convergence
RP/0/RP0/CPU0:router(config-rcmd)#protocol OSPF
RP/0/RP0/CPU0:router(config-rcmd-proto)#priority high
RP/0/RP0/CPU0:router(config-rcmd-proto-prio)#leaf-network 100
RP/0/RP0/CPU0:router(config-rcmd-proto-prio)#threshold 1
```

Related Commands

| Command | Description |
|--------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------|
| router-convergence, on page 1452 | 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 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.

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/RP0/CPU0:router (config) #router
RP/0/RP0/CPU0:router (config) #router-convergence
RP/0/RP0/CPU0:router (config-rcmd) #protocol OSPF
RP/0/RP0/CPU0:router (config-rcmd-proto) #priority high
RP/0/RP0/CPU0:router (config-rcmd-proto-prio) #leaf-network 100
RP/0/RP0/CPU0:router (config-rcmd-proto-prio) #threshold 1
```

| Related Commands | Command | Description |
|------------------|--------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------|
| | router-convergence, on page 1452 | Configures route convergence monitoring and enters router convergence monitoring and diagnostics (rcmd) configuration mode. |

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 | Parameter | Description |
|--------------------|-------------------------------|------------------------------------------------------------------------------|
| | diagnostics | Specifies an absolute directory path for storing diagnostic reports. |
| | <i>directory-path</i> | Specifies the path of the absolute directory for storing diagnostic reports. |
| | diagnostics-size | Specifies the maximum size of diagnostics directory. |
| | <i>maximum-directory-size</i> | Specified the size of the diagnostics directory. The range is 5% to 80%. |
| | reports | Specifies an absolute directory path for storing reports. |
| | <i>directory-path</i> | 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 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.

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/RP0/CPU0:router#configure
RP/0/RP0/CPU0:router (config) #router-convergence
RP/0/RP0/CPU0:router (config-rcmd) #storage-location
RP/0/RP0/CPU0:router (config-rcmd-store) #diagnostics /harddisk:/rcmd_logs
RP/0/RP0/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 1452 | Configures route convergence monitoring and enters router convergence monitoring and diagnostics (rcmd) configuration mode. |