

Routing Command Reference for Cisco 8000 Series Routers

First Published: 2020-03-01

Last Modified: 2026-09-08

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

© 2022 Cisco Systems, Inc. All rights reserved.



CONTENTS

PREFACE

Preface xvii

Preface xvii

Communications, Services, and Additional Information xvii

CHAPTER 1

BFD Commands 1

bfd 2

bfd address-family ipv4 destination 4

bfd address-family ipv4 echo minimum-interval 5

bfd address-family ipv4 fast-detect 6

bfd address-family ipv4 minimum-interval 7

bfd address-family ipv4 multiplier 9

bfd address-family ipv4 timers 12

bundle coexistence bob-blb 13

bfd dampening 15

bfd fast-detect 18

bfd minimum-interval 20

bfd mode 22

bfd multipath include location 23

bfd multiplier 24

clear bfd counters 26

clear bfd dampening 28

echo disable 30

interface (BFD) 32

show bfd 34

show bfd client 37

show bfd counters 38

show bfd hw-offload voq-usage 40
show bfd mib session 42
show bfd multipath 45
show bfd session 47
show bfd summary 53
show bundle (BFD) 55

CHAPTER 2 IS-IS Commands 57
address-family (IS-IS) 60
address-family multicast topology (IS-IS) 61
advertise route-policy 62
authentication-check disable 63
auto-cost (IS-IS) 64
circuit-type 66
clear isis process 68
clear isis route 69
clear isis statistics 71
csnp-interval 72
default-information originate (IS-IS) 73
fast-reroute per-link (IS-IS) 75
fast-reroute per-prefix (IS-IS) 77
fast-reroute per-link priority-limit (IS-IS) 79
fast-reroute per-prefix load-sharing disable (IS-IS) 80
fast-reroute per-prefix tiebreaker (IS-IS) 81
hello-interval (IS-IS) 82
hello-multiplier 84
hello-padding 86
hello-password 88
hello-password keychain 90
interface (IS-IS) 91
is-type 92
iid disable 94
log adjacency changes (IS-IS) 95
log pdu drops 96

lsp-gen-interval 97
lsp-mtu 99
lsp-fast-flooding 100
lsp-password 102
lsp-refresh-interval 104
maximum-paths (IS-IS) 105
maximum-redistributed-prefixes (IS-IS) 106
maximum-propagated-prefixes (IS-IS) 107
max-lsp-lifetime 109
max-lsp-tx 110
metric (IS-IS) 111
metric fallback anomaly 113
metric-style wide 115
mpls ldp auto-config 117
mpls ldp sync (IS-IS) 118
mpls traffic-eng (IS-IS) 119
mpls traffic-eng router-id (IS-IS) 120
oor-set-overload-bit disable 121
net 122
nsf (IS-IS) 124
passive (IS-IS) 125
point-to-point 126
protocol shutdown (IS-IS) 127
psnp-interval 128
redistribute (IS-IS) 129
remote-psnp-delay 133
router isis 135
router isis static neighbor 136
set default-isis-instance 137
set-overload-bit 139
show isis 141
show isis adjacency 143
show isis adjacency-log 145
show isis checkpoint adjacency 147

show isis checkpoint interface 149
show isis checkpoint lsp 150
show isis database 152
show isis database-log 161
show isis fast-reroute 163
show isis hostname 166
show isis interface 168
show isis lsp-log 172
show isis mesh-group 175
show isis mpls traffic-eng adjacency-log 177
show isis mpls traffic-eng advertisements 179
show isis mpls traffic-eng tunnel 181
show isis neighbors 183
show isis protocol 186
show isis purge-log 191
show isis route 193
show isis spf-log 197
show isis statistics 204
show isis topology 208
show protocols (IS-IS) 211
single-topology 214
snmp-server traps isis 215
spf-interval 217
spf-interval ietf 219
spf prefix-priority (IS-IS) 221
summary-prefix (IS-IS) 223

CHAPTER 3 **OSPF Commands** 225
address-family (OSPF) 228
adjacency stagger 229
area (OSPF) 231
authentication (OSPF) 232
authentication-key (OSPF) 234
auto-cost (OSPF) 236

capability opaque disable 237
clear ospf process 238
clear ospf redistribution 239
clear ospf routes 240
clear ospf statistics 241
cost (OSPF) 243
cost-fallback (OSPF) 245
database-filter all out (OSPF) 246
dead-interval (OSPF) 247
default-cost (OSPF) 249
default-information originate (OSPF) 250
default-metric (OSPF) 252
demand-circuit (OSPF) 253
disable-dn-bit-check 255
distance (OSPF) 256
distance ospf 258
distribute-list 259
domain-id (OSPF) 261
domain-tag 263
exchange-timer 264
fast-reroute (OSPFv2) 266
fast-reroute per-link exclude interface 268
fast-reroute per-prefix exclude interface (OSPFv2) 269
fast-reroute per-prefix lfa-candidate (OSPFv2) 270
fast-reroute per-prefix remote-lfa (OSPFv2) 271
fast-reroute per-prefix ti-lfa 272
fast-reroute per-prefix use-candidate-only (OSPFv2) 273
flood-reduction (OSPF) 274
hello-interval (OSPF) 275
ignore lsa mospf 277
interface (OSPF) 278
log adjacency changes (OSPF) 280
loopback stub-network 281
max external-lsa 282

max-lsa 284
max-metric 286
maximum interfaces (OSPF) 289
maximum paths (OSPF) 290
maximum redistributed-prefixes (OSPF) 291
message-digest-key 293
mpls ldp auto-config (OSPF) 296
mpls ldp sync (OSPF) 297
mpls traffic-eng (OSPF) 298
mpls traffic-eng igp-intact (OSPF) 300
mpls traffic-eng multicast-intact (OSPF) 301
mpls traffic-eng ldp-sync-update (OSPF) 302
mpls traffic-eng router-id (OSPF) 303
mtu-ignore (OSPF) 305
multi-area-interface 307
neighbor (OSPF) 309
neighbor database-filter all out 311
network (OSPF) 312
nsf (OSPF) 314
nsf flush-delay-time (OSPF) 316
nsf interval (OSPF) 317
nsf lifetime (OSPF) 318
nsr (OSPF) 319
nssa (OSPF) 320
ospf name-lookup 321
packet-size (OSPF) 322
passive (OSPF) 324
priority (OSPF) 326
protocol shutdown 328
queue dispatch flush-lsa 329
queue dispatch incoming 330
queue dispatch rate-limited-lsa 331
queue dispatch spf-lsa-limit 332
queue limit 333

range (OSPF) 334
redistribute (OSPF) 336
retransmit-interval (OSPF) 340
route-policy (OSPF) 342
router-id (OSPF) 343
router ospf (OSPF) 345
security ttl (OSPF) 346
segment-routing sr-prefer prefix-list 348
sham-link 350
show ospf 351
show ospf border-routers 355
show ospf database 357
show ospf flood-list 372
show ospf interface 374
show ospf mpls traffic-eng 376
show ospf message-queue 381
show ospf neighbor 383
show ospf request-list 390
show ospf retransmission-list 392
show ospf routes 394
show ospf sham-links 399
show ospf summary-prefix 401
show ospf virtual-links 403
show protocols (OSPF) 405
snmp context (OSPF) 407
snmp trap (OSPF) 409
snmp trap rate-limit (OSPF) 410
spf prefix-priority (OSPFv2) 411
stub (OSPF) 413
summary-prefix (OSPF) 414
timers lsa group-pacing 416
timers lsa min-arrival 417
timers throttle lsa all (OSPF) 418
timers throttle spf (OSPF) 420

transmit-delay (OSPF) 422
virtual-link (OSPF) 424
vrf (OSPF) 426

CHAPTER 4 **OSPFv3 Commands** 427

address-family (OSPFv3) 429
area (OSPFv3) 430
authentication (OSPFv3) 431
auto-cost (OSPFv3) 433
capability vrf-lite (OSPFv3) 434
clear ospfv3 process 435
clear ospfv3 redistribution 436
clear ospfv3 routes 437
clear ospfv3 statistics 438
cost (OSPFv3) 440
database-filter all out (OSPFv3) 442
dead-interval (OSPFv3) 443
default-cost (OSPFv3) 445
default-information originate (OSPFv3) 446
default-metric (OSPFv3) 448
demand-circuit (OSPFv3) 449
distance ospfv3 451
distribute-list prefix-list in 452
distribute-list prefix-list out 454
domain-id (OSPFv3) 456
encryption 457
flood-reduction (OSPFv3) 459
graceful-restart (OSPFv3) 460
hello-interval (OSPFv3) 461
instance (OSPFv3) 463
interface (OSPFv3) 465
log adjacency changes (OSPFv3) 467
maximum interfaces (OSPFv3) 468
maximum paths (OSPFv3) 469

maximum redistributed-prefixes (OSPFv3) 470
mtu-ignore (OSPFv3) 472
neighbor (OSPFv3) 474
network (OSPFv3) 476
nssa (OSPFv3) 478
nsr (OSPFv3) 480
ospfv3 name-lookup 481
packet-size (OSPFv3) 482
passive (OSPFv3) 483
priority (OSPFv3) 485
protocol shutdown (OSPFv3) 486
range (OSPFv3) 487
redistribute (OSPFv3) 489
retransmit-interval (OSPFv3) 493
router-id (OSPFv3) 495
router ospfv3 497
sham-link (OSPFv3) 498
show ospfv3 499
show ospfv3 border-routers 505
show ospfv3 database 507
show ospfv3 flood-list 519
show ospfv3 interface 521
show ospfv3 message-queue 524
show ospfv3 neighbor 526
show ospfv3 request-list 532
show ospfv3 retransmission-list 534
show ospfv3 routes 536
show ospfv3 statistics rib-thread 538
show ospfv3 summary-prefix 540
show ospfv3 virtual-links 542
show protocols (OSPFv3) 544
snmp context (OSPFv3) 546
snmp trap (OSPFv3) 548
snmp trap rate-limit (OSPFv3) 549

spf prefix-priority (OSPFv3) 550
stub (OSPFv3) 551
stub-router 552
summary-prefix (OSPFv3) 554
timers lsa arrival 556
timers pacing flood 557
timers pacing lsa-group 558
timers pacing retransmission 560
timers throttle lsa all (OSPFv3) 561
timers throttle spf (OSPFv3) 563
trace (OSPFv3) 565
transmit-delay (OSPFv3) 567
virtual-link (OSPFv3) 568
vrf (OSPFv3) 569

CHAPTER 5 **RIB Commands** 571
address-family next-hop dampening disable 573
clear route 574
maximum prefix (RIB) 576
lcc 577
rcc 578
recursion-depth-max 579
router rib 580
rump always-replicate 581
show lcc statistics 582
show rcc 584
show rcc statistics 586
show rcc vrf 588
show rib 589
show rib afi-all 591
show rib attributes 593
show rib client-id 594
show rib clients 596
show rib extcomms 598

show rib firsthop 600
show rib history 602
show rib next-hop 604
show rib opaques 607
show rib protocols 609
show rib recursion-depth-max 611
show rib statistics 613
show rib tables 615
show rib trace 617
show rib vpn-attributes 619
show rib vrf 621
show route 622
show route backup 629
show route best-local 632
show route connected 634
show route local 636
show route longer-prefixes 638
show route next-hop 640
show route quarantined 642
show route resolving-next-hop 644
show route static 646
show route summary 648
table ip-only activate vrf 651
tag-map tag 652

CHAPTER 6 RPL Commands 653
as-path in 655
as-path-set 657
community matches-any 659
community matches-every 661
community-set 663
destination in 666
done 668
drop 670

drop (ePBR) 671
end-policy 672
end-set 673
if 675
local-preference 680
med 681
next-hop-type is 682
pass 683
prefix-set 684
rib-has-route 686
route-policy (RPL) 687
set community 689
set extcommunity transitive-bandwidth 691
set local-preference 692
set med 693
set next-hop 695
set origin 697
show rpl 698
show rpl active as-path-set 700
show rpl active community-set 702
show rpl active extcommunity-set 704
show rpl active prefix-set 707
show rpl active rd-set 709
show rpl active route-policy 711
show rpl as-path-set 713
show rpl as-path-set attachpoints 714
show rpl as-path-set references 716
show rpl community-set 718
show rpl community-set attachpoints 720
show rpl community-set references 722
show rpl extcommunity-set 724
show rpl inactive as-path-set 727
show rpl inactive community-set 729
show rpl inactive extcommunity-set 731

show rpl inactive prefix-set 733
show rpl inactive rd-set 735
show rpl inactive route-policy 737
show rpl maximum 739
show rpl policy-global references 741
show rpl prefix-set 743
show rpl prefix-set attachpoints 744
show rpl prefix-set references 746
show rpl rd-set 748
show rpl rd-set attachpoints 749
show rpl rd-set references 751
show rpl route-policy 753
show rpl route-policy attachpoints 756
show rpl route-policy inline 758
show rpl route-policy references 760
show rpl route-policy uses 763
show rpl unused as-path-set 766
show rpl unused community-set 769
show rpl unused extcommunity-set 772
show rpl unused prefix-set 773
show rpl unused rd-set 776
show rpl unused route-policy 777
transmit (ePBR) 780

CHAPTER 7 **Static Routing** 781

address-family (static) 782
address-family ipv4 unicast (static) 783
address-family ipv6 unicast (static) 784
maximum path (static) 785
router static address-family ipv4 unicast 786
route (static) 787
router static 790
show static interfaces 791
show static vrf 792

vrf (static) 793



Preface

This preface contains these sections:

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

Preface

Initial release of the cumulative command reference document that covers all updates from Release 7.0.0 onwards.

Communications, Services, and Additional Information

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

Cisco Bug Search Tool

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



BFD Commands

- bfd, on page 2
- bfd address-family ipv4 destination, on page 4
- bfd address-family ipv4 echo minimum-interval, on page 5
- bfd address-family ipv4 fast-detect, on page 6
- bfd address-family ipv4 minimum-interval, on page 7
- bfd address-family ipv4 multiplier, on page 9
- bfd address-family ipv4 timers, on page 12
- bundle coexistence bob-blb, on page 13
- bfd dampening, on page 15
- bfd fast-detect, on page 18
- bfd minimum-interval, on page 20
- bfd mode, on page 22
- bfd multipath include location, on page 23
- bfd multiplier, on page 24
- clear bfd counters, on page 26
- clear bfd dampening, on page 28
- echo disable, on page 30
- interface (BFD), on page 32
- show bfd, on page 34
- show bfd client, on page 37
- show bfd counters, on page 38
- show bfd hw-offload voq-usage, on page 40
- show bfd mib session, on page 42
- show bfd multipath, on page 45
- show bfd session, on page 47
- show bfd summary, on page 53
- show bundle (BFD) , on page 55

bfd

To enter Bidirectional Forwarding Detection (BFD) configuration mode, use the **bfd** command in mode. To exit BFD configuration mode and return to global configuration mode, use the **no** form of this command.

bfd
no bfd

Command Default

No default behavior or values

Command Modes

Usage Guidelines

When you issue the **bfd** command in 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:

```
Router(config)# bfd
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 Router mode
  show       Show contents of configuration
```

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 configuration mode:

```
Router # configure
Router(config)# bfd
Router(config-bfd)#
```

Related Commands

Command	Description
Echo disable	
Interface bfd	
Show bfd	

bfd address-family ipv4 destination

To specify the destination address for BFD sessions on bundle member links, use the **bfd address-family ipv4 destination** command in interface configuration mode. To return to the default, use the **no** form of this command.

bfd address-family ipv4 destination *ip-address*
no bfd address-family ipv4 destination *ip-address*

Syntax Description

ip-address	32-bit IPv4 address in dotted-decimal format (A.B.C.D).

Command Default

No destination IPv4 address is configured.

Command Modes

Interface configuration (config-if)

Usage Guidelines

This command is supported on bundle interfaces only.

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:

```
Router# configure
Router(config)# interface Bundle-Ether 1
Router(config-if)# bfd address-family ipv4 destination 10.20.20.1
```

bfd address-family ipv4 echo minimum-interval

To specify the minimum interval for echo packets on IPv4 BFD sessions on bundle member links, use the **bfd address-family ipv4 echo minimum-interval** command in interface configuration mode. To return to the default, use the **no** form of this command.

Note The interface configuration of the echo timer value takes precedence over the global echo timer value.

bfd address-family ipv4 echo minimum-interval *milliseconds*

Syntax Description

milliseconds	Shortest interval between sending BFD echo packets to a neighbor. The range is 15 to 2000 milliseconds.

Command Default

The default value is the product of the async and multiplier values.

Command Modes

Interface configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

This command is supported on only on BFD over Bundle Per-Member Link Mode using Cisco standard.

This command allows to configure an echo interval to a value greater than async x bfd multiplier value.

Task ID

Task ID	Operations
bundle	read, write

Examples

The following example specifies that echo packets will be sent at a minimum interval of 900 ms for IPv4 BFD sessions on member links of an Ethernet bundle:

```
Router# configure
Router(config)# interface Bundle-Ether 1
Router(config-if)# bfd address-family ipv4 minimum-interval 200
Router(config-if)# bfd address-family ipv4 echo minimum-interval 900
```

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

Command Default

BFD sessions are disabled.

Command Modes

Interface configuration (config-if)

Usage Guidelines

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:

```
Router# configure
Router(config)# interface Bundle-Ether 1
Router(config-if)# bfd address-family ipv4 fast-detect
```

Related Commands

Command	Description
Bfd address family ipv4 destination	

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.

Command Default

The default is 150 ms.

Command Modes

Interface configuration (config-if)

Usage Guidelines

This command is supported on bundle interfaces only.

The BFD minimum interval is used with a configurable multiplier (**bfd address-family ipv4 multiplier** command) to determine the intervals and failure detection times for both control and echo packets in asynchronous mode on bundle member links.

For example, with a session interval of *I* and a multiplier of *M*, the following packet intervals and failure detection times apply for BFD asynchronous mode:

- Value of *I*—Minimum period between sending of BFD control packets.
- Value of *I* x *M*
  - BFD control packet failure detection time. This is the maximum amount of time that can elapse without receipt of a BFD control packet before the session is declared down.
  - Minimum period between sending of BFD echo packets.
- Value of (*I* x *M*) x *M*—BFD echo packet failure detection time. This is the maximum amount of time that can elapse without receipt of a BFD echo packet before the session is declared down.

When used with bundled VLANs, the following restrictions apply:

- The command specifies control packet intervals only because echo packets are not supported.
- The minimum interval is 250 ms.

The **bfd address-family ipv4 minimum-interval** command in bundle interface configuration overrides the minimum intervals specified by the **bfd minimum-interval** command in other areas of BFD configuration.

Note When multiple applications share the same BFD session, the application with the most aggressive timer is used locally. Then, the result is negotiated with the peer router.

Keep the following router-specific rules in mind when configuring the minimum BFD interval:

- The maximum rate in packets-per-second (pps) for BFD sessions is linecard-dependent. If you have multiple linecards supporting BFD, then the maximum rate for BFD sessions per system is the supported linecard rate multiplied by the number of linecards.
- The maximum number of all BFD sessions on the router is 1024.
- The maximum number of all BFD sessions on the router is 1440.

To calculate the rate for BFD sessions on bundle members running in asynchronous mode without echo:

- Divide 1000 by the value of the minimum interval (as specified by the **bfd address-family ipv4 minimum-interval** command). This is also the base rate used per member session with echo:

  Asynchronous rate per bundle member = (1000 / Min-interval)

To calculate the rate for BFD sessions on bundle members running in asynchronous mode with echo:

- Determine the echo interval, which is the value of the minimum interval (specified by the **bfd address-family ipv4 minimum-interval** command) multiplied by the multiplier value (specified by the **bfd address-family ipv4 multiplier** command).

  Echo interval = (Min-interval x Multiplier)

- Calculate the overall rate supported for all members on the bundle:

  Ethernet bundle rate = (1000 / Echo interval) x 64

- Add the asynchronous base rate per bundle member to find the total rate for all bundle links:

  Total bundle rate = Ethernet bundle rate + (Base asynchronous rate x Number of links)

Task ID

Task ID	Operations
bundle	read, write

Examples

The following example specifies that control packets will be sent at a minimum interval of 200 ms for IPv4 BFD sessions on member links of an Ethernet bundle:

```
Router# configure
Router(config)# interface Bundle-Ether 1
Router(config-if)# bfd address-family ipv4 minimum-interval 200
```

Related Commands

Command	Description
Bfd minimum-interval	
Bfd address-family ipv4 multiplier	

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)

Usage Guidelines

This command is supported on bundle interfaces only.

The BFD multiplier is used with a configurable minimum interval (**bfd address-family ipv4 minimum-interval** command) to determine the intervals and failure detection times for both control and echo packets in asynchronous mode on bundle member links.

For example, with a session interval of *I* and a multiplier of *M*, the following packet intervals and failure detection times apply for BFD asynchronous mode:

- Value of *I*—Minimum period between sending of BFD control packets.
- Value of *I* x *M*
  - BFD control packet failure detection time. This is the maximum amount of time that can elapse without receipt of a BFD control packet before the session is declared down.
  - Minimum period between sending of BFD echo packets.



Note The maximum echo packet interval for BFD on bundle member links is the minimum of either 30 seconds or the asynchronous control packet failure detection time.

- Value of (*I* x *M*) x *M*—BFD echo packet failure detection time. This is the maximum amount of time that can elapse without receipt of a BFD echo packet before the session is declared down.

Keep the following router-specific rules in mind when configuring the minimum BFD interval:

- The maximum rate in packets-per-second (pps) for BFD sessions is linecard-dependent. If you have multiple linecards supporting BFD, then the maximum rate for BFD sessions per system is the supported linecard rate multiplied by the number of linecards.
- The maximum number of all BFD sessions per linecard is 1024.
- The maximum number of all BFD sessions per linecard is 1440.

To calculate the rate for BFD sessions on bundle members running in asynchronous mode without echo:

- Divide 1000 by the value of the minimum interval (as specified by the **bfd address-family ipv4 minimum-interval** command). This is also the base rate used per member session with echo:

  Asynchronous rate per bundle member = (1000 / Min-interval)

To calculate the rate for BFD sessions on bundle members running in asynchronous mode with echo:

- Determine the echo interval, which is the value of the minimum interval (specified by the **bfd address-family ipv4 minimum-interval** command) multiplied by the multiplier value (specified by the **bfd address-family ipv4 multiplier** command).

  Echo interval = (Min-interval x Multiplier)
- Calculate the overall rate supported for all members on the bundle:

  Ethernet bundle rate = (1000 / Echo interval) x 64
- Add the asynchronous base rate per bundle member to find the total rate for all bundle links:

  Total bundle rate = Ethernet bundle rate + (Base asynchronous rate x Number of links)

Task ID

Task ID	Operations
bundle	read, write

Task ID Examples

The following example specifies the following packet intervals and failure detection times for IPv4 BFD sessions on member links with asynchronous echo mode on an Ethernet bundle:

- 200 ms control packet interval
- 600 ms control packet failure detection interval
- 600 ms echo packet interval
- 1800 ms echo packet failure detection interval

```
Router# configure
Router(config)# interface Bundle-Ether 1
Router(config-if)# bfd address-family ipv4 minimum-interval 200
Router(config-if)# bfd address-family ipv4 multiplier 3
```

Related Commands

Command	Description
Bfd address family ipv4 minimum-interval	
Bfd minimum-interval	

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)

Usage Guidelines

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:

```
Router# configure
Router(config)# interface Bundle-Ether 1
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:

```
Router# configure
Router(config)# interface Bundle-Ether 1
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

Usage Guidelines

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":

```
Router#configure
Router(config)#bfd
Router(config-bfd)#bundle coexistence bob-blb inherit
```

This example shows how to configure the coexitence mechanism between BFD over Bundle (BoB) and BFD over Logical Bundle (BLB) as being "logical":

```
Router#configure
Router(config)#bfd
Router(config-bfd)#bundle coexistence bob-blb logical
```

bfd dampening

To configure a device to dampen a flapping Bidirectional Forwarding Detection (BFD) session, use the **bfd dampening** command in global configuration mode. To restore the values of the dampening timers to their default values, use the **no** form of this command.

bfd dampening [**bundle-member**] { **initial-wait** | **l3-only-mode** | **maximum-wait** | **secondary-wait** | **extension down-monitoring** | **threshold** } *milliseconds*
no bfd dampening [**bundle-member**] { **initial-wait** | **l3-only-mode** | **maximum-wait** | **secondary-wait** | **extension down-monitoring** | **threshold** } *milliseconds*

Syntax Description

bundle-member	Specifies initial, maximum, or secondary delays in milliseconds for BFD session startup on BFD bundle members.
initial-wait	Specifies the initial delay in milliseconds before starting a BFD session. For bundle members, the default is 16000. For non-bundle interfaces, the default is 2000.
l3-only-mode	Applies dampening, if the failure is Layer 3 (L3)-specific.
maximum-wait	Specifies the maximum delay in milliseconds before starting a BFD session. **Note** The maximum delay must be greater than the initial delay.
secondary-wait	Specifies a secondary delay in milliseconds before starting a BFD session. For bundle members, the default is 20000. For non-bundle interfaces, the default is 5000.
extensions down-monitoring	Enables session monitoring extensions in down state.
threshold	Specifies a stability threshold to enable dampening in milliseconds. Range 6000 to 360000. The default is 120000.
milliseconds	For bundle members, the range is 1-518400000. For non-bundle interfaces, the range is 1-3600000.

Command Default BFD dampening is enabled by default.

Command Modes Global configuration

Command History

Release	Modification
Release 3.9.0	This command was introduced.
Release 4.0.0	The **bundle-member** keyword was added.
Release 4.2.0	The **l3-only-mode** keyword was added to the other dampening options.
Release 5.1	The **extensions** and **down-monitoring** keywords were added.

Usage Guidelines

Note The **initial-wait**, **maximum-wait**, and **secondary-wait** keywords are dampening timers.

You do not have to configure the BFD startup timers. If you do configure the initial wait startup timer (using the **initial-wait** keyword), then it must be less than the value of the maximum-wait timer.

You do not have to configure the BFD startup timers. If you do configure the initial wait startup timer (using the **initial-wait** keyword), then it must be less than the value of the maximum-wait timer.

By default, BFD dampening is applied to all sessions in the following manner:

- If a session is brought down, then dampening is applied before a session is allowed to transition to initial/up states.
- Length of time a session is dampened grows exponentially with continuous session flap.
- If a session remains up for minimum two minutes, then the length of time a session dampens with the next session flap is reset to the initial dampening value.

BFD on bundle member applies dampening, only if the detected failure is specific to layer 3. BFD dampening is not invoked for L1 or L2 failures. BFD is started after Layer 1 and Layer 2 (LACP) is up to prevent a race condition and false triggers. BFD is notified to stop/ignore when L1 or L2 goes down and must be notified to start/resume when L1 or L2 recovers for a given/affected link/member.

BFD applies dampening till the session transitions from up to down state and the session is not removed. Whenever there is a failure detected at L1 or L2 , the bundle manager removes BFD session on a member.

When dampening is removed a syslog message ‘Exponential backoff dampening for BFD session has been cleared for specified BFD session. When/if same session gets created by application(s), only calculated initial wait time will be applied’ is generated. If this is the desired behavior, then dampening can be enabled by configuring the BFD configuration, by using the command **bfd dampening bundle-member l3-failure-only**.

This example shows how to configure BFD dampening by specifying an initial and maximum delay for BFD session startup on BFD bundle members:

```
Router# configure
Router(config)# bfd dampening bundle-member initial-wait 1000
```

```
Router(config)# bfd dampening bundle-member maximum-wait 3000
```

This example shows how to configure BFD dampening on a non-bundle interface:

```
Router# configure
Router(config)# bfd dampening initial-wait 1000
Router(config)# bfd dampening maximum-wait 3000
```

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.
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 6.1x	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.

In OSPF and OSPF3 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.

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, you must enter the **no bfd fast-detect** command.

Task ID

Task ID	Operations
bgp	read, write
isis	read, write
ospf	read, write

Examples

The following example shows how to configure BFD on a BGP router:

```
Router(config)# route-policy pass-all
Router(config-rpl)# pass
Router(config-rpl)# end-policy
Router(config-rpl)# end
Router# configure
Router(config)# router bgp 65000
Router(config-bgp)# bfd minimum-interval 3
Router(config-bgp)# address-family ipv4 unicast\
Router(config-bgp-af)# exit
Router(config-bgp)# neighbor 192.168.70.24
Router(config-bgp-nbr)# remote-as 2
Router(config-bgp-nbr)# address-family ipv4 unicast
Router(config-bgp-nbr)# bfd fast-detect
Router(config-bgp-nbr)# route-policy pass-all in
Router(config-bgp-nbr)# route-policy pass-all out
```

bfd minimum-interval

To specify the minimum control packet interval for BFD sessions for the corresponding BFD configuration scope, use the **bfd minimum-interval** command in the appropriate configuration mode. To return the router to the default setting, use the **no** form of this command.

bfd minimum-interval *milliseconds*
no bfd minimum-interval [*milliseconds*]

Syntax Description

milliseconds	Interval between sending BFD hello packets to the neighbor. The range is 300 to 30000 milliseconds.

Command Modes

Router configuration

Interface configuration

Router configuration

Area configuration

Area interface configuration

Interface configuration

Command History

Release	Modification
Release 7.3.1	This command was introduced.

Usage Guidelines

In OSPF environments, the setting of the **bfd minimum-interval** command is inherited from the highest-level configuration mode in which the command was configured. From the lowest to the highest configuration modes, the inheritance rules are as follows:

- If you configure the minimum interval in area interface configuration mode, the updated interval affects the BFD sessions on the specified interface only.
- If you configure the minimum interval in area configuration mode, the updated interval affects the BFD sessions on all interfaces in the specified area.

If 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.

- If a session is running in asynchronous mode without echo, then PPS used for this session is (1000 / asynchronous interval in milliseconds).
- If a session is running in asynchronous mode with echo, then PPS used for this session is (1000 / echo interval in milliseconds).

  This is calculated as: 1000 / value of the **bfd minimum-interval** command.

- The table below defines the maximum number of BFD sessions that is allowed per line card.

Table 1: Supported Scale

Timer Value	Maximum Number of Sessions Supported in a line card.	Maximum Number of Sessions Supported in a system.
300 ms	250	500
1 second	250	500
2 seconds	250	500

Examples

The following example shows how to set the BFD minimum interval for a BGP routing process:

```
Router(config)# router bgp 6500
Router(config-bgp)# bfd minimum-interval 300
```

bfd mode

To enable the option to use Cisco or IETF mode for BFD over bundle, use the **bfd mode** command in interface configuration mode. To disable the option to use Cisco or IETF mode for BFD over bundle, use the **no** form of this command.

bfd mode { **cisco** | **ietf** }

no bfd mode

Syntax Description

cisco	Specifies the use of Cisco mode for BFD over bundle.
ietf	Specifies the use of IETF mode for BFD over bundle.

Command Default

The default member mode is **cisco**.

Command Modes

Interface configuration (config-if)

Usage Guidelines

If the BFD mode is configured when the bundle is being created, the configuration goes through. This is because, both the BFD state as well as the bundle state are 'down' during bundle creation. To apply the mode change for existing sessions, bring down and then recreate the BFD sessions for that bundle.

A BFD session on the member interfaces can belong to only one mode (Cisco or IETF mode). Mix of the modes within the same bundle is not supported. This command is supported on bundle interfaces only.

Task ID

Task ID	Operations
bundle	read, write

Examples

This example shows how to enable **ietf** mode for the BFD session on an Ethernet bundle interface:

```
Router# configure
Router(config)#  interface bundle-ether 1
Router(config-if)#   bfd mode ietf
```

bfd multipath include location

To include specific linecards to host BFD multiple path sessions, use the **bfd multipath include location** command in the XR configuration mode. To remove the configuration, use the **no** form of this command.

bfd multipath include location *node-id*
no bfd multipath include location *node-id*

Syntax Description

location *node-id*	Configures BFD multipath on the specified location. The *node-id* variable is mentioned in the *rack/slot/module* notation.

Command Default

No default behavior or values

Command Modes

XR Config mode

Command History

Release	Modification
Release 6.3x	This command was introduced.

Task ID

Task ID	Operation
bgp	read, write
ospf	read, write
isis	read, write

Example

This example shows how to run the **bfd multipath include location** command on a specific location:

```
Router# configure
Router(config)# bfd multipath include location 0/5/CPU0
```

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

Command Default

The default multiplier is 3.

Command Modes

Router configuration

Interface configuration

Router configuration

Area configuration

Area interface configuration

Interface configuration

Usage Guidelines

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
ospf	read, write

Examples

The following example shows how to set the BFD multiplier in a BGP routing process:

```
Router(config)# router bgp 65000
Router(config-bgp)# bfd multiplier 3
```

clear bfd counters

To clear Bidirectional Forwarding Detection (BFD) counters, use the **clear bfd counters** command in mode.

clear bfd counters { **ipv4** | [**singlehop** |] | [**singlehop** |] | **all** | **label** } [**packet**] [**timing**] [**interface** *type* *interface-path-id*] **location** *node-id*

Syntax Description

ipv4	(Optional) Clears BFD over IPv4 information only.
all	(Optional) Clears BFD over IPv4 information.
packet	(Optional) Specifies that packet counters are cleared.
timing	(Optional) Specifies that timing counters are cleared.
interface	(Optional) Specifies the interface from which the BFD packet counters are cleared.
type	Specifies the interface type. For more information, use the question mark (**?**) online help function.
interface-path-id	Physical interface or virtual interface. **Note** Use the **show interfaces** command to see a list of all interfaces currently configured on the router. For more information about the syntax for the router, use the question mark (**?**) online help function.
location *node-id*	Clears BFD counters from the specified location. The *node-id* argument is entered in the *rack/slot/module* notation.

Command Default

The default is the default address family identifier (AFI) that is set by the **set default-afi** command.

Command Modes

Usage Guidelines

For the *interface-path-id* argument, use the following guidelines:

- If specifying a physical interface, the naming notation is *rack*/*slot*/*module*/*port*. The slash between values is required as part of the notation. An explanation of each component of the naming notation is as follows:
  - *rack*: Chassis number of the rack.
  - *slot*: Physical slot number of the line card.
  - *module*: Module number. A physical layer interface module (PLIM) is always 0.
  - *port*: Physical port number of the interface.
- If specifying a virtual interface, the number range varies, depending on interface type.

Examples

The following example shows how to clear the BFD IPv4 timing counters:

```
Router# clear bfd counters ipv4 timing location 0/5/cpu0
```

Related Commands

Command	Description
Bfd fast detect	
Bfd minimum-interval	
Bfd multiplier	
Show bfd	

clear bfd dampening

To reset BFD dampening counters, use the **clear bfd dampening** command in mode.

clear bfd dampening { [**interface** *type interface-path-id*] | **location** *node-id all* | **type** [**bundle-logical** | **bundle-per-member**] }

Syntax Description

interface	Specifies the interface from which the BFD dampening sessions are cleared.
type	Specifies the interface type. For more information, use the question mark (**?**) online help function.
interface-path-id	Physical interface or virtual interface. **Note** Use the **show interfaces** command to see a list of all interfaces currently configured on the router. For more information about the syntax for the router, use the question mark (**?**) online help function.
location *node-id*	Clears BFD dampening sessions from the specified location. The *node-id* argument is entered in the *rack/slot/module* notation.
all	(Optional) Clears BFD dampening session on all locations.
type	Specifies the BFD session type. For more information, use the question mark (**?**) online help function.
bundle-per-member	Clears BFD dampening over the member links of BFD over bundle sessions.
bundle-logical	Clears BFD dampening of the BFD over logical bundle session.

Command Modes

Usage Guidelines

For the *interface-path-id* argument, use the following guidelines:

- If specifying a physical interface, the naming notation is *rack*/*slot*/*module*/*port*. The slash between values is required as part of the notation. An explanation of each component of the naming notation is as follows:
  - *rack*: Chassis number of the rack.
  - *slot*: Physical slot number of the line card.
  - *module*: Module number. A physical layer interface module (PLIM) is always 0.
  - *port*: Physical port number of the interface.
- If specifying a virtual interface, the number range varies, depending on interface type.

Examples

The following example shows how to clear BFD dampening:

- To clear BFD dampening on all BFD sessions:

```
Router# clear bfd dampening location all
```

- To clear BFD dampening on BFD sessions over bundle interfaces:

```
Router# clear bfd dampening type bundle-per-member
```

- To clear bfd dampening on BFD sessions over bundle logical interfaces:

```
Router# clear bfd dampening type bundle-logical
```

- To clear bfd dampening on BFD sessions over specific interface:

```
Router# clear bfd dampening interface gigabitEthernet 0/3/0/0
```

clear bfd dampening log sample

The following is a sample syslog output of the **clear bfd dampening** command:

```
LC/0/3/CPU0:Oct  6 16:52:40.632 : bfd_agent[127]: %L2-BFD-6-SESSION_DAMPENING_CLR : Dampening
 for BFD session to neighbor 10.1.1.2
on interface GigabitEthernet0/3/0/0 has been cleared
```

echo disable

To disable echo mode on a router or on an individual interface or bundle, use the **echo disable** command in Bidirectional Forwarding Detection (BFD) configuration mode. To return the router to the default configuration where echo mode is enabled, use the **no** form of this command.

echo disable
no echo disable

Command Default

No default behavior or values

Command Modes

BFD configuration

BFD interface configuration

Usage Guidelines

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:

```
Router# configure
Router(config)# bfd
Router(config-bfd)# echo disable
```

The following example shows how to disable echo mode on an individual interface:

```
Router# configure
Router(config)# bfd
Router(config-bfd)# interface gigabitethernet 0/1/0/0
Router(config-bfd-if)# echo disable
```

Related Commands

Command	Description
Bfd	Enters BFD configuration mode.
Interface bfd	Enters BFD interface configuration mode.
ipv4 verify unicast source reachable-via	Enables and disables IPv4 uRPF checking on an IPv4 interface.
Show bfd	Displays BFD information.

interface (BFD)

To enter Bidirectional Forwarding Detection (BFD) interface configuration mode, where you can disable echo mode on an interface, use the **interface** command in BFD configuration mode. To return to BFD configuration mode, use the **no** form of this command.

interface *type* *interface-path-id*
no **interface** *type* *interface-path-id*

Syntax Description

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

Command Default

No default behavior or values

Command Modes

BFD configuration

Usage Guidelines

For the *interface-path-id* argument, use the following guidelines:

- If specifying a physical interface, the naming notation is *rack*/*slot*/*module*/*port*. The slash between values is required as part of the notation. An explanation of each component of the naming notation is as follows:
  - *rack*: Chassis number of the rack.
  - *slot*: Physical slot number of the line card.
  - *module*: Module number. A physical layer interface module (PLIM) is always 0.
  - *port*: Physical port number of the interface.
- If specifying a virtual interface, the number range varies, depending on interface type.

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.

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:

```
Router# configure
Router(config)# bfd
Router(config-bfd)# interface gigabitethernet 0/1/0/0
Router(config-bfd-if)#
```

The following example shows how to enter BFD interface configuration mode for a Packet-over-SONET/SDH (POS) interface:

```
Router# configure
Router(config)# bfd
Router(config-bfd)# interface pos 0/1/0/0
Router(config-bfd-if)#
```

Related Commands

Command	Description
BFD	Enters BFD configuration mode.
Echo disable	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.
Show BFD	Displays BFD information.

show bfd

To display Bidirectional Forwarding Detection (BFD) information for a specific location, use the **show bfd** command in EXEC mode.

show bfd [**ipv4** | [**singlehop | multihop**] | **ipv6 [singlehop | multihop] | all** | **label**] **interface** [**destination | source**] [**location** *node-id*]

Syntax Description

ipv4	(Optional) Displays BFD over IPv4 information only.
singlehop	(Optional) Displays BFD singlehop information only.
ipv6	(Optional) Displays BFD over IPv6 information only.
all	(Optional) Displays both BFD over IPv4 and BFD over IPv6 information.
label	(Optional) Displays the BFD label information.
interface	Specifies the BFD interface.
destination	(Optional) Specifies the destination IPv4 unicast address.
source	(Optional) Specifies the source IPv4 unicast address.
location *node-id*	Displays BFD information for the specified location. The *node-id* argument is entered in the *rack/slot/module* notation.

Command Default

The default is the default address family identifier (AFI) that is set by the **set default-afi** command.

Command Modes

EXEC mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

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** command:

```
Router# show bfd
```

```
IPV4 Sessions Up: 0, Down: 0, Total: 0
```

The following example shows the output from the **show bfd all** command:

```
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:

```
Router# show bfd ipv4

IPV4 Sessions Up: 0, Down: 0, Total: 0
```

The following example shows the output from the **show bfd ipv4 location** command:

```
Router# show bfd ipv6 session detail interface tenGigE 0/0/0/0.100

I/f: TenGigE0/0/0/0.100, Location: 0/0/CPU0
Dest: 1001:1:1:1:1:1:1:2
Src: 1001:1:1:1:1:1:1:1
 State: UP for 0d:1h:37m:3s, number of times UP: 1
 Session type: PR/V6/SH
Received parameters:
 Version: 1, desired tx interval: 300 ms, required rx interval: 300 ms
 Required echo rx interval: 0 us, multiplier: 3, diag: None
 My discr: 2148076695, your discr: 2148075550, state UP, D/F/P/C/A: 0/0/0/1/0
Transmitted parameters:
 Version: 1, desired tx interval: 300 ms, required rx interval: 300 ms
 Required echo rx interval: 0 us, multiplier: 3, diag: None
 My discr: 2148075550, your discr: 2148076695, state UP, D/F/P/C/A: 0/1/0/1/0
Timer Values:
 Local negotiated async tx interval: 300 ms
 Remote negotiated async tx interval: 300 ms
 Desired echo tx interval: 0 s, local negotiated echo tx interval: 0 us
 Echo detection time: 0 us(0 us*3), async detection time: 900 ms(300 ms*3)
Local Stats:
 Intervals between async packets:
   Tx: Number of intervals=3, min=10 ms, max=2290 ms, avg=824 ms
       Last packet transmitted 5823 s ago
   Rx: Number of intervals=6, min=3 ms, max=2 s, avg=461 ms
       Last packet received 5823 s ago
Intervals between echo packets:
   Tx: Number of intervals=0, min=0 s, max=0 s, avg=0 s
       Last packet transmitted 0 s ago
   Rx: Number of intervals=0, min=0 s, max=0 s, avg=0 s
       Last packet received 0 s ago
 Latency of echo packets (time between tx and rx):
   Number of packets: 0, min=0 us, max=0 us, avg=0 us
Session owner information:
```

```
                          Desired              Adjusted
  Client                Interval   Multiplier Interval   Multiplier
  --------------------- ---------- ---------- ---------- ----------
  ipv6_static           300 ms     3          300 ms     3
  ipv6_static           300 ms     3          300 ms     3

H/W Offload Info:
 H/W Offload capability : Y, Hosted NPU     : 0/0/CPU0/NPU0
 Async Offloaded        : Y, Echo Offloaded : N
 Async rx/tx            : 7/4

Platform Info:
NPU ID: 0
Async RTC ID        : 1          Echo RTC ID        : 0
Async Feature Mask  : 0x8        Echo Feature Mask  : 0x0
Async Session ID    : 0x3c07     Echo Session ID    : 0x0
Async Tx Key        : 0x3c070801  Echo Tx Key        : 0x0
Async Tx Stats addr : 0x3f69e800   Echo Tx Stats addr : 0x0
Async Rx Stats addr : 0x4069e800   Echo Rx Stats addr : 0x0
```

show bfd client

To display Bidirectional Forwarding Detection (BFD) client information, use the **show bfd client** command in XR EXEC mode 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 mode

Release	Modification
Release 7.0.12	This feature was introduced.

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 2: show bfd client Field Descriptions

Field	Description
Name	Name of the BFD client.
Node	Location of the BFD client.
Num sessions	Number of active sessions for the BFD client.

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** | **mulithop**] | **ipv6** [**singlehop** | **mulithop**] **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 multihopinformation only.
all	(Optional) Displays BFD over IPv4 and BFD over IPv6 information.
packet	Specifies that packet counters are displayed.
interface	(Optional) Specifies the interface for which to show counters.
type	Interface type. For more information, use the question mark (**?**) online help function.
interface-path-id	Physical interface or virtual interface. **Note** Use the **show interfaces** command to see a list of all interfaces currently configured on the router. For more information about the syntax for the router, use the question mark (**?**) online help function.
location *node-id*	Displays BFD counters from the specified location. The *node-id* argument is entered in the *rack/slot/module* notation.

Command Default

The default is the default address family identifier (AFI) that is set by the **set default-afi** command.

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

For the *interface-path-id* argument, use the following guidelines:

- If specifying a physical interface, the naming notation is *rack*/*slot*/*module*/*port*. The slash between values is required as part of the notation. An explanation of each component of the naming notation is as follows:
  - *rack*: Chassis number of the rack.
  - *slot*: Physical slot number of the line card.

- *module*: Module number. A physical layer interface module (PLIM) is always 0.
- *port*: Physical port number of the interface.

- If specifying a virtual interface, the number range varies, depending on interface type.

Task ID

Task ID	Operations
bgp	read
ospf	read
isis	read
mpls-te	read

Examples

The following example displays the output from the **show bfd counters packet** for both IPv4.

```
RP/0/RP0/CPU0:router# show bfd counters ipv4 packet

IPv4 Singlehop:
 GigabitEthernet0/0/1/2                  Recv       Xmit                 Recv       Xmit
      Async:                             4148       4137       Echo: (   47136)     80192
 GigabitEthernet0/1/1/2                  Recv       Xmit                 Recv       Xmit
      Async:                             116876     125756       Echo: ( 2268192)   2301312
 Bundle-Ether10                          Recv       Xmit                 Recv       Xmit
      Async:                             2          0       Echo:        0          0
 Bundle-Ether20                          Recv       Xmit                 Recv       Xmit
      Async:                             91         0       Echo:         0          0

IPv4 Multihop: (Src IP/Dst IP/Vrf Id)
 10.15.151.4/10.16.151.4/0x12345678      Recv       Xmit
      Async:                             0       570337
```

show bfd hw-offload voq-usage

To display VOQ assignments for hardware-offloaded BFD sessions, use the **show bfd hw-offload voq-usage** command in EXEC mode.

show bfd hw-offload voq-usage { **all** | **discriminator** *discriminator* } **location** *location-id*

Syntax Description

all	Displays every eligible hardware-offloaded BFD session.
discriminator *discriminator*	Displays the hardware-offloaded BFD session that has the specified local discriminator.
location *location-id*	Displays VOQ information for the specified location. Enter the *location-id* argument in *rack/slot/module* notation.

Command Default

None

Command Modes

EXEC mode

Command History

Release	Modification
Release 26.3.1	This command was introduced.

Usage Guidelines

- Enter the local discriminator that a BFD session command displays to get the detailed output.
- The detailed output presents the discriminator in decimal and hexadecimal formats.
- The command supports hardware-offloaded IPv4 and IPv6 single-hop BFD sessions.
- The command does not display multihop BFD sessions.

Note For BFD over Logical Bundle (BLB), BFD over Bundle (BOB), and tunnel sessions, the command can display None for fields that have no corresponding session value.

Task ID

Task ID	Operations
bgp	read
ospf	read
isis	read
mpls-te	read

Examples

This example displays the VOQ assignments for all eligible sessions:

```
Router# show bfd hw-offload voq-usage all location 0/RP0/CPU0

0/RP0/CPU0
-----------------
Discr        Interface    Type       State      TX VOQ     RX VOQ     NPU      Slot   SysPort
-----        ---------    ----       -----      ------     ------     ---      ----   -------
2147549185   Hu0/0/0/2    SH-v4      UP         656        288        0        0      16
2147549189   Hu0/1/0/1    SH-v6      UP         680        288        0        1      13

Total BFD HW-offloaded sessions: 2
```

This example displays detailed VOQ information for one session:

```
Router# show bfd hw-offload voq-usage discriminator 2147549189 location 0/RP0/CPU0

0/RP0/CPU0
-----------------
Session Discriminator: 2147549189 (0x80010005)
Interface:             Hu0/1/0/1 (0x780000cc)
Session Type:          SH-v6
Local State:           UP
TX VOQ Base:           680
RX VOQ Base:           288
NPU ID:                0
Slot:                  1
System Port:           13
```

Table 3: Output field descriptions for show bfd hw-offload voq-usage

Field	Description
`Discr`	Identifies the local session discriminator.
`Interface`	Identifies the egress interface for the session.
`Type`	Identifies the BFD session type.
`State`	Displays the local BFD session state.
`TX VOQ`	Displays the transmit VOQ base queue number.
`RX VOQ`	Displays the receive punt VOQ base queue number.
`NPU`	Identifies the network processing unit.
`Slot`	Identifies the line card slot.
`SysPort`	Identifies the system port.
`Total BFD HW-offloaded sessions`	Displays the total number of hardware-offloaded BFD sessions.

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 *node-id*	(Optional) Displays all IPv4 and IPv6 BFD MIB session information stored on the specified node. The *node-id* argument is entered in the *rack/slot/module* notation.

Command Default

When *node-id* is not specified, information for all IPv4 and IPv6 BFD MIB sessions, stored on the route processor node, is displayed.

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

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:

```
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:

```
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 4: show bfd mib Field Descriptions

Field	Description
date and timestamp	Date and time stamp during which a snapshot of the BFD MIB session information is taken.
Local Discr	Local discriminator (in decimal and hexadecimal) that uniquely identifies the BFD MIB session.
Remote Discr	Session discriminator (in decimal and hexadecimal) that was chosen by the remote system for the BFD MIB session.
BFD session	Index of interface upon which the BFD MIB session is running. Also, neighboring IP address that is monitored with the BFD MIB session.
Current State	Current state of the BFD MIB session.
Number of Times UP	Number of times the BFD MIB session has gone into the up state since the router was last rebooted.
Running Version	BFD protocol version number in which the BFD MIB session is running.
Last Down Diag	Diagnostic value associated with the last time the BFD MIB session went down.
Last Up Time (s.ns)	Value of sysUpTime, in *seconds.nanoseconds*, during which the BFD MIB session last came up. If such an event does not exist, a zero is displayed.
Last Down Time (s.ns)	Value of sysUpTime, in *seconds.nanoseconds*, during which communication was last lost with the neighbor. If such an event does not exist, a zero is displayed.
Detection Multiplier	Failure detection multiplier.

Field	Description
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.

show bfd multipath

To display information concerning only BFD multipath sessions, use the **show bfd multipath** command in the XR EXEC mode mode.

show bfd multipath {**ipv4** | **ipv6** | **label** | **all**} **location** *node-id*

Syntax Description

ipv4	Displays BFD over IPv4 information only.
ipv6	Displays BFD over IPv6 information only.
label	Displays BFD label information.
all	Displays both BFD over IPv4 and BFD over IPv6 information.
location *node-id*	Displays BFD counters from the specified location. The *node-id* argument is entered in the *rack/slot/module* notation.

Command Default

No default behavior or values

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

Task ID

Task ID	Operation
bgp	read
ospf	read
isis	read
mpls-te	read

Example

This example shows the sample output for **show bfd multipath** command:

```
RP/0/RP0/CPU0:router#show bfd multipath location 0/5/cpu0

Int/Src Addr    Label/Dest Addr VRF ID     Discr      Node      State
--------------- --------------- ---------- ---------- --------- --------
```

```
pw-ether 1      10.10.10.10     0x00000002 0x4        0/5/CPU0  DOWN
tunnel-ip 1     1.1.1.1         0x8        0x5        0/5/CPU0  UP
```

show bfd session

To display Bidirectional Forwarding Detection (BFD) session information, use the **show bfd session** command in XR Exec mode.

show bfd [**ipv4** | [**singlehop** | **multihop**] | [**singlehop** | **multihop**] | **all**] **session** [**interface** *type interface-path-id* [**destination** *ip-address*] [**detail**]] **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.
interface	(Optional) Specifies the interface for which to show information.
type	Interface type. For more information, use the question mark (?) online help function.
interface-path-id	Physical interface or virtual interface. **Note** Use the **show interfaces** command to see a list of all interfaces currently configured on the router. For more information about the syntax for the router, use the question mark (**?**) online help function.
destination *ip-address*	(Optional) Displays the BFD session destined for the specified IP address.
detail	(Optional) Displays detailed session information, including statistics and number of state transitions.

location *node-id*	(Optional) Displays BFD sessions hosted from the specified location. The *node-id* argument is entered in the *rack/slot/module* notation.

Command Default

The default is the default address family identifier (AFI) that is set by the **set default-afi** command.

Command Modes

XR EXEC mode

Usage Guidelines

For the *interface-path-id* argument, use the following guidelines:

- If specifying a physical interface, the naming notation is *rack*/*slot*/*module*/*port*. The slash between values is required as part of the notation. An explanation of each component of the naming notation is as follows:
  - *rack*: Chassis number of the rack.
  - *slot*: Physical slot number of the line card.
  - *module*: Module number. A physical layer interface module (PLIM) is always 0.
  - *port*: Physical port number of the interface.
- If specifying a virtual interface, the number range varies, depending on interface type.

Examples

The following example shows the output from the **show bfd session** command with the **detail** keyword and IPv4 as the default:

```
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
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
```

```
Session owner information:
  Client           Desired interval         Multiplier
  ---------------- --------------------     --------------
  bgp-             250 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:

```
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
```

Table 5: 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 *rack/slot/module* notation
dest	IP address of the destination endpoint.
src	IP address of the source endpoint.
State	Current state of the connection, and the number of days, hours, minutes, and seconds that this connection has been active.
number of times UP	Number of times this connection has been brought up.

Field	Description
Received parameters	Provides information on the last transmitted control packet for the session: • Version—Version number of the BFD protocol. • desired tx interval—Desired transmit interval. • required rx interval—Required receive interval. • Required echo rx interval—Required echo receive interval. • multiplier— Number of times a packets is missed before BFD declares the neighbor down. • diag—diagnostic code specifying the peer system's reason for the last transition of the session from Up to some other state. • My discr—unique, nonzero discriminator value generated by the transmitting system, used to demultiplex multiple BFD sessions between the same pair of systems. • your discr— discriminator received from the corresponding remote system. This field reflects back the received value of My discr, or is zero if that value is unknown.
Transmitted parameters	Provides information on the last transmitted control packet for the session: • Version—Version number of the BFD protocol. • desired tx interval—Desired transmit interval. • required rx interval—Required receive interval • Required echo rx interval—Required echo receive interval • multiplierNumber of times a packets is missed before BFD declares the neighbor down. • diag—diagnostic code specifying the local system's reason for the last transition of the session from Up to some other state. • My discr—unique, nonzero discriminator value generated by the transmitting system, used to demultiplex multiple BFD sessions between the same pair of systems. • your discr— discriminator received from the corresponding remote system. This field reflects back the received value of My discr, or is zero if that value is unknown.

Field	Description
Timer Values	Provides information on the timer values used by the local and remote ends, as follows: • Local negotiated async tx interval—interval at which control packets are being transmitted by the local end. • Remote negotiated async tx interval—interval at which control packets should be transmitted by the remote end. • Desired echo tx interval—interval at which the local end would like to transmit echo packets. • local negotiated echo tx interval—interval at which echo packets are being transmitted by the local end. • Echo detection time—local failure detection time of echo packets. It is the product of the local negotiated echo tx interval and the local multiplier. • async detection time—local failure detection time of the asynchronous mode (control packets). It is the product of the remote negotiated async tx interval and the remote multiplier.
Local Stats	Displays the local transmit and receive statistics, • Intervals between async packets—provides measurements on intervals between control packets (tx and rx): • Number of intervals—number of sampled intervals between control packets • min—minimum measured interval between 2 consecutive control packets • max—maximum measured interval between 2 consecutive control packets • avg—average measured interval between 2 consecutive control packets • Last packet received/transmitted—indicates how long ago the last control packet was received/transmitted. • Intervals between echo packets—provides measurements on intervals between echo packets (tx and rx). The measurements have the same meaning as for async packets. • Latency of echo packets (time between tx and rx)—provides measurements on latency of echo packets, i.e. the time between tx and rx of echo packets: • Number of packets—number of sampled echo packets. • min—minimum measured latency for echo packets. • max—maximum measured latency of echo packets. • avg—average measured latency of echo packets.

Field	Description
Session owner information	Provides the following information about the session owner. • Client—name of the client application process. • Desired interval—desired interval provided by the client, in milliseconds. • Multiplier—multiplier value provided by the client.

Use the **show bfd label** session to verify the configuration at each BFD label.

```
Router#show bfd label session
Fri Dec  2 17:12:21.497 UTC

Interface           Label              Local det time(int*mult)      State
                                     Echo             Async
H/W                 NPU
------------------- --------------- ---------------- ---------------- ----------
tt1 (LSP:103)       24001           n/a              150ms(50ms*3)    UP
Yes                 0/1/CPU0
tt2 (LSP:102)       24002           n/a              150ms(50ms*3)    UP
Yes                 0/1/CPU0
tt3 (LSP:101)       24004           n/a              150ms(50ms*3)    UP
Yes                 0/1/CPU0
tt4 (LSP:103)       24005           n/a              150ms(50ms*3)    UP
Yes                 0/1/CPU0
tt5 (LSP:104)       24006           n/a              150ms(50ms*3)    UP
Yes                 0/1/CPU0
```

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 *node-id*	Displays BFD counters from the specified location. The *node-id* argument is entered in the *rack/slot/module* notation.

Command Default

No default behavior or values

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

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:

```
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
```

show bundle (BFD)

To display the IPV6 BFD configured, use the **show bundle** command in the EXEC mode.

Example

This example shows the output of the show bundle bundle-ether command with IPv6 Bfd configured:

```
#sh bundle bundle-ether 1

bundle-ether1
  Status:                                  Down
  Local links <active/standby/configured>:   0 / 0 / 1
  Local bandwidth <effective/available>:     0 (0) kbps
  MAC address (source):                    000a.000b.000c (Configured)
  Inter-chassis link:                      No
  Minimum active links / bandwidth:        1 / 1 kbps
  Maximum active links:                    64
  Wait while timer:                        2000 ms
  Load balancing:
    Link order signaling:                  Not configured
    Hash type:                             Default
    Locality threshold:                    None
  LACP:                                    Operational
    Flap suppression timer:                Off
    Cisco extensions:                      Disabled
    Non-revertive:                         Enabled
  mLACP:                                   Not configured
IPv4 BFD:                                 Not configured
  IPv6 BFD:                                Operational
    State:                                 Down
    Mode:                                  Not configured
    Fast detect:                           Enabled
    Start timer:                           Off
    Neighbor-unconfigured timer:           Off
    Preferred min interval:                150 ms
    Preferred multiple:                    3
    Destination address:                   1::1
```



IS-IS Commands

- address-family (IS-IS), on page 60
- address-family multicast topology (IS-IS), on page 61
- advertise route-policy, on page 62
- authentication-check disable, on page 63
- auto-cost (IS-IS), on page 64
- circuit-type, on page 66
- clear isis process, on page 68
- clear isis route, on page 69
- clear isis statistics, on page 71
- csnp-interval, on page 72
- default-information originate (IS-IS), on page 73
- fast-reroute per-link (IS-IS), on page 75
- fast-reroute per-prefix (IS-IS), on page 77
- fast-reroute per-link priority-limit (IS-IS), on page 79
- fast-reroute per-prefix load-sharing disable (IS-IS), on page 80
- fast-reroute per-prefix tiebreaker (IS-IS), on page 81
- hello-interval (IS-IS), on page 82
- hello-multiplier, on page 84
- hello-padding, on page 86
- hello-password, on page 88
- hello-password keychain, on page 90
- interface (IS-IS), on page 91
- is-type, on page 92
- iid disable, on page 94
- log adjacency changes (IS-IS), on page 95
- log pdu drops, on page 96
- lsp-gen-interval, on page 97
- lsp-mtu, on page 99
- lsp-fast-flooding, on page 100
- lsp-password, on page 102
- lsp-refresh-interval, on page 104
- maximum-paths (IS-IS), on page 105
- maximum-redistributed-prefixes (IS-IS), on page 106

- maximum-propagated-prefixes (IS-IS), on page 107
- max-lsp-lifetime, on page 109
- max-lsp-tx, on page 110
- metric (IS-IS), on page 111
- metric fallback anomaly, on page 113
- metric-style wide, on page 115
- mpls ldp auto-config , on page 117
- mpls ldp sync (IS-IS), on page 118
- mpls traffic-eng (IS-IS), on page 119
- mpls traffic-eng router-id (IS-IS), on page 120
- oor-set-overload-bit disable, on page 121
- net, on page 122
- nsf (IS-IS), on page 124
- passive (IS-IS), on page 125
- point-to-point, on page 126
- protocol shutdown (IS-IS), on page 127
- psnp-interval, on page 128
- redistribute (IS-IS), on page 129
- remote-psnp-delay, on page 133
- router isis, on page 135
- router isis static neighbor, on page 136
- set default-isis-instance, on page 137
- set-overload-bit, on page 139
- show isis, on page 141
- show isis adjacency, on page 143
- show isis adjacency-log, on page 145
- show isis checkpoint adjacency, on page 147
- show isis checkpoint interface, on page 149
- show isis checkpoint lsp, on page 150
- show isis database, on page 152
- show isis database-log, on page 161
- show isis fast-reroute, on page 163
- show isis hostname, on page 166
- show isis interface, on page 168
- show isis lsp-log, on page 172
- show isis mesh-group, on page 175
- show isis mpls traffic-eng adjacency-log, on page 177
- show isis mpls traffic-eng advertisements, on page 179
- show isis mpls traffic-eng tunnel, on page 181
- show isis neighbors, on page 183
- show isis protocol, on page 186
- show isis purge-log, on page 191
- show isis route, on page 193
- show isis spf-log, on page 197
- show isis statistics, on page 204
- show isis topology, on page 208

- show protocols (IS-IS), on page 211
- single-topology, on page 214
- snmp-server traps isis, on page 215
- spf-interval, on page 217
- spf-interval ietf, on page 219
- spf prefix-priority (IS-IS), on page 221
- summary-prefix (IS-IS), on page 223

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 XR Config mode or interface configuration mode. To disable support for an address family, use the **no** form of this command.

address-family { **ipv4** | **ipv6** } { **unicast** | **multicast** }

Syntax Description

ipv4	Specifies IPv4 address prefixes.
ipv6	Specifies IPv6 address prefixes.
unicast	Specifies unicast address prefixes.
multicast	Specifies multicast address prefixes.

Command Default

An address family is not specified. The default subaddress family (SAFI) is unicast.

Command Modes

XR Config mode

Interface configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

Use the **address family** command to place the router or interface in address family configuration mode. In router address family configuration mode, you can configure routing that uses standard IPv4 or IPv6 address prefixes. An address family must be specified in interface configuration mode. In interface address family configuration mode, you can alter interface parameters for IPv4or IPv6.

You must specify an address family in order to configure parameters that pertain to a single address family.

Task ID

Task ID	Operations
isis	read, write

Examples

The following example shows how to configure the IS-IS router process with IPv4 unicast address prefixes:

```
RP/0/RP0/CPU0:router(config)# router isis isp
RP/0/RP0/CPU0:router(config-isis)# interface HundredGigE 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)#
```

address-family multicast topology (IS-IS)

To enable a multicast topology when configuring Intermediate System-to-Intermediate System (IS-IS) routing (or to place a given topology within the IS-IS interface), use the **address-family multicast topology** command with either IPv4 or IPv6 address prefix in the appropriate configuration mode. To disable a multicast topology in IS-IS, use the **no** form of this command.

address-family { **ipv4** | **ipv6** } [**maximum prefix** *prefix-limit*]
no address-family

Syntax Description

ipv4	Specifies IPv4 address prefixes.
ipv6	Specifies IPv6 address prefixes.
topology	Specifies the name of the topology.
maximum prefix	Specifies maximum number of prefixes that a routing table can have.
prefix-limit	Maximum number of prefixes. Range is from 32 to 2,000,000.

Command Default

An address family for multicast topology is not specified. The default subaddress family (SAFI) is unicast.

Command Modes

Router configuration

Interface configuration

Task ID

Task ID	Operations
isis	read, write

advertise route-policy

To specify the route policy that sets the conditions under which the static neighbor is advertised, use the **advertise** command on the IS-IS static neighbor configuration mode.

advertise route-policy *route-policy*

Syntax Description

route-policy	Name of the configured policy.

Command Default

No route policy is applied for advertising the static neighbor.

Command Modes

IS-IS static neighbor configuration

Command History

Release	Modification
Release 26.1.1	The command was introduced.

Usage Guidelines

When **passive** command is enabled on the interface with the static neighbor, the static adjacency and the local IP subnet on the interface are advertised.

When **suppressed** command is configured on the interface with the static neighbor, advertising of the local subnet stops, but the static adjacency continues to be advertised.

When **shutdown** command is configured on the interface with the static adjacency, all advertising, including the static adjacency, stops.

Task ID

Task ID	Operation
isis	read and write

Example

This example shows how to use a route policy to control the advertisement of the static neighbor.

```
Router(config)# router isis 1
Router(config-isis)# affinity-map RED bit-position 7
Router(config-isis)# interface GigabitEthernet0/2/0/7
Router(config-isis-if)# static-neighbor
Router(config-isis-static-nbr)# system-id 0001.0002.0007
Router(config-isis-static-nbr)# remote ipv4 address 10.1.1.2
Router(config-isis-static-nbr)# advertise route-policy isis-bgp
Router(config-isis-static-nbr)# exit
Router(config-isis-static)# address-family ipv4 unicast
Router(config-if)# exit
Router(config-if)# address-family ipv6 unicast
Router(config-if)# exit
```

authentication-check disable

To suppress Intermediate System-to-Intermediate System (IS-IS) authentication check, use the **authentication-check disable** command in configuration mode. To remove this function, use the **no** form of this command.

authentication-check disable
no authentication-check disable

Command Default

Authentication check is enabled

Command Modes

XR Config mode

Command History

Release	Modification
Release 7.7.1	This command was introduced.

Task ID

Task ID	Operations
isis	read, write

Examples

The following example shows how to disable authentication check for IS-IS.

```
Router# configure
Router(config)# router isis 1
Router(config)# authentication-check disable
Router(config)# commit
```

auto-cost (IS-IS)

To control how the Intermediate System-to-Intermediate System (IS-IS) 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** *bandwidth* [**granularity** *granularity*]

Syntax Description

reference-bandwidth *bandwidth*	Specifies a bandwidth in kilobits per second (kbps). The range is from 1 to 4294967295.
granularity *granularity*	Specifies a granularity in kilobits per second (kbps). The range is from 1 to 4294967295.

Command Default

mbps : 100 Mbps

Command Modes

Router configuration

Command History

Release	Modification
Release 25.4.1	This command was introduced.

Usage Guidelines

The interface-specific **address-family auto-cost reference-bandwidth** configuration overrides the router-level IS-IS **address-family auto-cost reference-bandwidth** configuration.

The router-level IS-IS **address-family auto-cost reference-bandwidth** configuration applies to all interfaces except those with an interface-specific **address-family auto-cost reference-bandwidth** configuration.

Task ID

Task ID	Operations
isis	read, write

Examples

The following example shows how to set the reference value for the auto cost calculation to 30000000 kbps:

```
Router# configure
Router# (config)#router isis 1
Router(config-isis)# address-family ipv4 unicast
Router# (config-isis-af)#auto-cost reference-bandwidth 30000000 granularity 100000
```

The following example shows how to set the reference value for auto-cost calculation on an interface to 30,000,000 kbps.:

```
Router# configure
  Router# (config)#router isis 1
```

```
Router(config-isis)# interface Bundle-Ether1
Router(config-isis-if)# address-family ipv4 unicast
Router(config-if-if-af)# auto-cost reference-bandwidth 20000000 granularity 100000
```

circuit-type

To configure the type of adjacency used for the Intermediate System-to-Intermediate System (IS-IS) protocol, use the **circuit-type** command in interface configuration mode. To reset the circuit type to Level l and Level 2, use the **no** form of this command.

circuit-type { **level-1** | **level-1-2** | **level-2-only** }
no circuit-type

Syntax Description

level-1	Establishes only Level 1 adjacencies over an interface.
level-1-2	Establishes both Level 1 and Level 2 adjacencies, if possible.
level-2-only	Establishes only Level 2 adjacencies over an interface.

Command Default

Default adjacency types are Level 1 and Level 2 adjacencies.

Command Modes

Interface configuration

Usage Guidelines

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 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	Configures the routing level for an instance of the IS-IS routing process.
Net	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 mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

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
```

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** | **multicast** | **safi-all**} [**topology** *topo-name*] **route**

Syntax Description

instance *instance-id*	(Optional) Specifies IS-IS sessions for the specified IS-IS instance only. • The *instance-id* argument is the instance identifier (alphanumeric) defined by the **router isis** command.
afi-all	Specifies IP Version 4 (IPv4) and IP Version 6 (IPv6) address prefixes.
ipv4	Specifies IPv4 address prefixes.
ipv6	Specifies IPv6 address prefixes.
unicast	Specifies unicast address prefixes.
multicast	Specifies multicast address prefixes.
safi-all	Specifies all secondary address prefixes.
topology *topo-name*	(Optional) Specifies topology table information and name of the topology table.

Command Default

No default behavior or value

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

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
```

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

instance *instance-id*	(Optional) Clears IS-IS sessions for the specified IS-IS instance only. • The *instance-id* argument is the instance identifier (alphanumeric) defined by the **router isis** command.
type	Interface type. For more information, use the question mark (?) online help function.
interface-path-id	Physical interface or virtual interface. **Note** Use the **show interfaces** command to see a list of all interfaces currently configured on the router. For more information about the syntax for the router, use the question mark (?) online help function.

Command Default

No default behavior or values

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

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
```

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** }]

Syntax Description

seconds	Interval (in seconds) of time between transmission of CSNPs on multiaccess networks. This interval applies only for the designated router. Range is 0 to 65535 seconds.
level { **1** | **2** }	(Optional) Specifies the interval of time between transmission of CSNPs for Level 1 or Level 2 independently.

Command Default

seconds : 10 seconds

Both Level 1 and Level 2 are configured if no level is specified.

Command Modes

Interface configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

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 HundredGigE 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 [**external** | **route-policy** *route-policy-name*]

Syntax Description

external	(Optional) Enables default route to be originated as an external route.
route-policy	(Optional) Defines the conditions for the default route.
route-policy-name	(Optional) Name for the route policy.

Command Default

A default route is not generated into an IS-IS routing domain.

Command Modes

Address family configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

If a router configured with the **default-information originate** command has a route to 0.0.0.0 in the routing table, IS-IS originates an advertisement for 0.0.0.0 in its link-state packets (LSPs).

Without a route policy, the default is advertised only in Level 2 LSPs. For Level 1 routing, there is another process to find the default route, which is to look for the closest Level 1 and Level 2 router. The closest Level 1 and Level 2 router can be found by looking at the attached-bit (ATT) in Level 1 LSPs.

A route policy can be used for two purposes:

- To make the router generate the default route in its Level 1 LSPs.
- To advertise 0.0.0.0/0 conditionally.

Task ID

Task ID	Operations
isis	read, write

Examples

The following example shows how to generate a default external route into an IS-IS domain:

```
RP/0/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
```

fast-reroute per-link (IS-IS)

To enable IP fast reroute (IPFRR) loop-free alternate (LFA) prefix independent per-link computation, use the **fast-reroute per-link** command in interface address family configuration mode. To disable this feature, use the **no** form of this command.

fast-reroute per-link [**exclude interface** *type interface-path-id* | **level** { **1** | **2** } | **lfa-candidate interface** *type interface-path-id*]

Syntax Description

exclude	Specifies fast-reroute (FRR) loop-free alternate (LFA) computation exclusion information
level {1 | 2}	Configures FRR LFA computation for one level only.
lfa-candidate	Specifies FRR LFA computation candidate information
interface	Specifies an interface that needs to be either excluded from FRR LFA computation (when used with **exclude** keyword) or to be included to LFA candidate list in FRR LFA computation (when used with the **lfa-candidate** keyword).
type	Interface type. For more information, use the question mark (?) online help function.
interface-path-id	Physical interface or virtual interface. **Note** Use the show interfaces command to see a list of all interfaces currently configured on the router. For more information about the syntax for the router, use the question mark (?) online help function.

Command Default

IP fast-reroute LFA per-link computation is disabled.

Command Modes

Interface address family configuration

Command History

Release	Modification
Release 7.0.12	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-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 HundredGigE 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. To disable this feature, use the **no** form of this command.

fast-reroute per-prefix [**exclude interface** *type interface-path-id* | **level** { **1** | **2** } | **lfa-candidate interface** *type interface-path-id* | **remote-lfa** { **maximum-metric** *metric-value* | **tunnel mpls-ldp** } **prefix-list** *prefix-list-name* [**level** { **1** | **2** }]]

Syntax Description

exclude	Specifies fast-reroute (FRR) loop-free alternate (LFA) computation exclusion information
level {1 | 2}	Configures FRR LFA computation for one level only.
lfa-candidate	Specifies FRR LFA computation candidate information
interface	Specifies an interface that needs to be either excluded from FRR LFA computation (when used with **exclude** keyword) or to be included to LFA candidate list in FRR LFA computation (when used with the **lfa-candidate** keyword).
type	Interface type. For more information, use the question mark (?) online help function.
interface-path-id	Physical interface or virtual interface. **Note** Use the show interfaces command to see a list of all interfaces currently configured on the router. For more information about the syntax for the router, use the question mark (?) online help function.
remote-lfa	Enable remote LFA related configuration.
prefix-list *prefix-list-name*	Filter PQ node router ID based on prefix list.

Command Default

IP fast-reroute LFA per-prefix computation is disabled.

Command Modes

Interface address family configuration

Command History

Release	Modification
Release 7.0.12	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)# interface HundredGigE 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-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** }

Syntax Description

critical	Enables LFA omputation for critical priority prefixes only.
high	Enables LFA computation for for criticaland high priority prefixes.
medium	Enables LFA computation for for critical, high, and medium priority prefixes.
level {1|2}	Sets priority-limit for routing Level 1 or Level 2 independently.

Command Default

Fast-reroute per link priority limit LFA computation is disabled.

Command Modes

IPv4 unicast address family configuration

IPv6 unicast address family configuration

IPv4 multicast address family configuration

IPv6 multicast address family configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

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-sharing disable

Syntax Description

level {1|2}	Disables load-sharing for Level 1 or Level 2 independently.

Command Default

Load sharing is enabled.

Command Modes

IPv4 unicast address family configuration

IPv4 multicast address family configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

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**] **index** *index_number* **level {1 | 2}**

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.
index	Sets preference order among tie-breakers.
index_number	Value for the index. Range is 1-255.
level {1 | 2}	Configures tiebreaker for Level 1 or Level 2 independently.

Command Default

Tie-breaker for multiple backups is not configured.

Command Modes

IPv4 unicast address family configuration

IPv4 multicast address family configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

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
```

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** }]

Syntax Description

seconds	Integer value (in seconds) for the length of time between consecutive hello packets. By default, a value three times the hello interval *seconds* is advertised as the *hold time* in the hello packets sent. (That multiplier of three can be changed by using the **hello-multiplier** command.) With smaller hello intervals, topological changes are detected more quickly, but there is more routing traffic. Range is 1 to 65535 seconds.
level { **1** | **2** }	(Optional) Specifies the hello interval for Level 1 and Level 2 independently. For broadcast interfaces only.

Command Default

seconds : 10 seconds

Both Level 1 and Level 2 are configured if no level is specified.

Command Modes

Interface configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

The hello interval can be configured independently for Level 1 and Level 2, except on serial point-to-point interfaces. (Because only a single type of hello packet is sent on serial links, it is independent of Level 1 or Level 2.) Configuring Level 1 and Level 2 independently is used on LAN interfaces.

Note A shorter hello interval gives quicker convergence, but increases bandwidth and CPU usage. It might also add to instability in the network.

A slower hello interval saves bandwidth and CPU. Especially when used in combination with a higher hello multiplier, this strategy may increase overall network stability.

For point-to-point links, IS-IS sends only a single hello for Level 1 and Level 2, making the **level** keyword meaningless on point-to-point links. To modify hello parameters for a point-to-point interface, omit the **level** keyword.

Note Currently, a user can configure an aggressive hello-interval (lower than the default of 10 seconds for peer-to-peer session). But, if NSR or NSF is configured, the default hello interval has to be used so that the sessions do not run into the risk of flapping during switchover.

Using LAN adjacencies in high availability (HA) scenarios is not recommended, since there is no designated intermediate system (DIS) redundancy in the protocol and traffic will either drop or be rerouted temporarily during DIS re-election.

Task ID

Task ID	Operations
isis	read, write

Examples

The following example shows how to configure HundredGigE 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 HundredGigE 0/6/0/0
RP/0/RP0/CPU0:router(config-isis-if)# hello-interval 5 level 1
```

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** }]

Syntax Description

multiplier	Advertised hold time in IS-IS hello packets is set to the hello multiplier times the hello interval. Range is 3 to 1000. Neighbors declare an adjacency to this down router after not having received any IS-IS hello packets during the advertised hold time. The hold time (and thus the hello multiplier and the hello interval) can be set on an individual interface basis, and can be different between different networking devices in one area. Using a smaller hello multiplier gives faster convergence, but can result in more routing instability. Increase the hello multiplier to a larger value to help network stability when needed. Never configure a hello multiplier to a value lower than the default value of 3.
level { **1** | **2** }	(Optional) Specifies the hello multiplier independently for Level 1 or Level 2 adjacencies.

Command Default

multiplier : 3

Both Level 1 and Level 2 are configured if no level is specified.

Command Modes

Interface configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

The “holding time” carried in an IS-IS hello packet determines how long a neighbor waits for another hello packet before declaring the neighbor to be down. This time determines how quickly a failed link or neighbor is detected so that routes can be recalculated.

Use the **hello-multiplier** command in circumstances where hello packets are lost frequently and IS-IS adjacencies are failing unnecessarily. You can raise the hello multiplier and lower the hello interval (hello-interval (IS-IS), on page 82 command) correspondingly to make the hello protocol more reliable without increasing the time required to detect a link failure.

On point-to-point links, there is only one hello for both Level 1 and Level 2. Separate Level 1 and Level 2 hello packets are also sent over nonbroadcast multiaccess (NBMA) networks in multipoint mode, such as X.25, Frame Relay, and ATM.

Task ID

Task ID	Operations
isis	read, write

Examples

The following example shows how the network administrator wants to increase network stability by making sure an adjacency goes down only when many (ten) hello packets are missed. The total time to detect link failure is 60 seconds. This strategy ensures that the network remains stable, even when the link is fully congested.

```
RP/0/RP0/CPU0:router(config)# router isis isp
RP/0/RP0/CPU0:router(config-isis)# interface HundredGigE 0/2/0/1
RP/0/RP0/CPU0:router(config-isis-if)# hello-interval 6
RP/0/RP0/CPU0:router(config-isis-if)# hello-multiplier 10
```

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 and process configuration mode. To suppress padding, use the **no** form of this command.

hello-padding { **adaptive** | **always** | **disable** | **sometimes** }

[**level** { **1** | **2** }]

Syntax Description

adaptive	Enable hello-padding till neighbor confirms adjacency up.
always	Always enable hello-padding.
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 and IS-IS process configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.
Release 7.10.1	This command was extended to IS-IS process configuration mode.

Usage Guidelines

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 HundredGigE 0/2/0/1:

```
RP/0/RP0/CPU0:router(config)# router isis isp
RP/0/RP0/CPU0:router(config-isis)# interface HundredGigE 0/2/0/1
RP/0/RP0/CPU0:router(config-isis-if)# hello-padding disable
```

The following example shows how to suppress IS-IS hello padding at the process level:

```
Router (config)#router isis 100
Router (config-isis)#hello-padding disable
Router (config)#commit
```

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**]

Syntax Description

hmac-md5	(Optional) Specifies that the password use HMAC-MD5 authentication.
text	(Optional) Specifies that the password use clear text password authentication.
clear	(Optional) Specifies that the password be unencrypted.
encrypted	(Optional) Specifies that the password be encrypted using a two-way algorithm.
password	Authentication password you assign for an interface.
level { **1** | **2** }	(Optional) Specifies whether the password is for a Level 1 or a Level 2 protocol data unit (PDU).
send-only	(Optional) Specifies that the password applies only to protocol data units (PDUs) that are being sent and does not apply to PDUs that are being received.

Command Default

Both Level 1 and Level 2 are configured if no level is specified.

password: encrypted text

Command Modes

Interface configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

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 HundredGigE 0/2/0/3 interface:

```
RP/0/RP0/CPU0:router(config)# router isis isp
RP/0/RP0/CPU0:router(config-isis)# interface HundredGigE 0/2/0/3
RP/0/RP0/CPU0:router(config-isis-if)# hello-password hmac-md5 clear mypassword
```

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**]

Syntax Description

keychain	Keyword that specifies the keychain to be configured. An authentication password keychain is a sequence of keys that are collectively managed and used for authenticating a peer-to-peer group.
keychain-name	Specifies the name of the keychain.
level { **1** | **2** }	(Optional) Specifies whether the keychain is for a Level 1 or a Level 2 protocol data unit (PDU).
send-only	(Optional) Specifies that the keychain applies only to protocol data units (PDUs) that are being sent and does not apply to PDUs that are being received.

Command Default

Both Level 1 and Level 2 are configured if no level is specified.

password: encrypted text

Command Modes

Interface configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

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 HundredGigE:

```
RP/0/RP0/CPU0:router(config)# router isis isp
RP/0/RP0/CPU0:router(config-isis)# interface HundredGigE 0/1/0/0
RP/0/RP0/CPU0:router(config-isis-if)# hello-password keychain mykeychain level 1 send-only
```

interface (IS-IS)

To configure the Intermediate System-to-Intermediate System (IS-IS) protocol on an interface, use the **interface** command in XR Config mode. To disable IS-IS routing for interfaces, use the **no** form of this command.

interface *type* *interface-path-id*

Syntax Description

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

Command Default

No interfaces are specified.

Command Modes

XR Config mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

An address family must be established on the IS-IS interface before the interface is enabled for IS-IS protocol operation.

Task ID

Task ID	Operations
isis	read, write

Examples

The following example shows how to enable an IS-IS multitopology configuration for IPv4 on HundredGigE 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 HundredGigE 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 HundredGigE 0/3/0/0
RP/0/RP0/CPU0:router(config-if)# ipv4 address 2001::1/64
```

is-type

To configure the routing level for an Intermediate System-to-Intermediate System (IS-IS) area, use the **is-type** command in XR Config 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** }

Syntax Description

level-1	Specifies that the router perform only Level 1 (intra-area) routing. This router learns only about destinations inside its area. Level 2 (interarea) routing is performed by the closest Level 1-2 router.
level-1-2	Specifies that the router perform both Level 1 and Level 2 routing.
level-2-only	Specifies that the routing process acts as a Level 2 (interarea) router only. This router is part of the backbone, and does not communicate with Level 1-only routers in its own area.

Command Default

Both Level 1 and Level 2 are configured if no level is specified.

Command Modes

XR Config mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

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
```

iid disable

To suppress instance-ID TLV for `hello` and `lsp` packets when multiple IS-IS instances are configured in a router, use the **iid disable** command in process configuration mode.

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

iid disable

Syntax Description

This command has no keywords or arguments.

Command Default

Sending or receiving is enabled.

Command Modes

IS-IS process configuration

Command History

Release	Modification
Release 7.10.1	The command was introduced.

Usage Guidelines

Each IS-IS instance has a unique instance-ID set, the TLV of which is sent in the **hello** and **lsp** packets.

Task ID

Task ID	Operations
isis	read, write

Examples

The following example shows how to suppress instance ID TLV at the IS-IS process configuration level:

```
Router(config-isis)#iid disable
Router(config-isis)#commit
```

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

Command Default

No IS-IS instance log messages are generated.

Command Modes

Router configuration

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

log pdu drops

Command Default

PDU logging is disabled.

Command Modes

XR Config mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

Use the **log pdu drops** command to monitor a network when IS-IS PDUs are suspected of being dropped. The reason for the PDU being dropped and current PDU drop statistics are recorded.

The following are examples of PDU logging output:

```
%ISIS-4-ERR_LSP_INPUT_Q_OVERFLOW - An incoming LSP or SNP pdu was dropped because the input
 queue was full
%ISIS-3-ERR_SEND_PAK - The process encountered a software-error while sending the IS-IS
packet
```

Task ID

Task ID	Operations
isis	read, write

Examples

The following example shows how to enable PDU logging:

```
RP/0/RP0/CPU0:router(config)# router isis isp
RP/0/RP0/CPU0:router(config-isis)# log pdu drops
```

lsp-gen-interval

To customize IS-IS throttling of link-state packet (LSP) generation, use the **lsp-gen-interval** command in XR Config 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** }]

Syntax Description

initial-wait *initial*	Specifies the initial LSP generation delay (in milliseconds). Range is 0 to 120000 milliseconds.
secondary-wait *secondary*	Specifies the hold time between the first and second LSP generation (in milliseconds). Range is 1 to 120000 milliseconds.
maximum-wait *maximum*	Specifies the maximum interval (in milliseconds) between two consecutive occurrences of an LSP being generated. Range is 1 to 120000 milliseconds.
level { **1** | **2** }	(Optional) Specifies the LSP time interval for Level 1 or Level 2 independently.

Command Default

initial-wait *initial* : 50 milliseconds

secondary-wait *secondary* : 200 milliseconds

maximum-wait *maximum* : 5000 milliseconds

Command Modes

XR Config mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

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
```

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 XR Config mode. To restore the default, use the **no** form of this command.

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 from 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

XR Config mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

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 168 command to display the value.

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
```

lsp-fast-flooding

To enable the fast flooding of Link State Packets (LSPs) to improve network response to topology changes, use the **lsp-fast-flooding** command in IS-IS router configuration mode.

lsp-fast-flooding [**max-lsp-tx** *lsps-per-sec* | **remote-psnp-delay** *milliseconds*\

Syntax Description

max-lsp-tx *lsps-per-sec*	(Optional) Sets the maximum number of LSPs that can be transmitted per second. The default is 1000. The range is from 33 through 5000
remote-psnp-delay *milliseconds*	(Optional) Specifies the maximum delay, in milliseconds, that the router should wait for a PSNP acknowledgment from neighbors after sending an LSP. The default is 500 ms. The range is from 1 through 5000 ms.

Command Default

This command is disabled by default. IS-IS fast flooding of LSPs is not enabled.

Command Modes

Router configuration

Command History

Release	Modification
Release 24.3.1	This command was introduced.

Usage Guidelines

None.

Task ID

Task ID	Operation
isis	read, write

Example

This example shows how to enable IS-IS fast flooding of LSPs.

```
Router# configure terminal
Router(config)# router isis 1
Router(config-router)# lsp-fast-flooding
```

This example shows how to set the maximum LSP transmission rate to 600 LSPs per second.

```
Router# configure terminal
Router(config)# router isis 1
Router(config-router)# lsp-fast-flooding
Router(config-isis-lsp-fast-flooding)# max-lsp-tx 600
```

This example shows how to set the PSNP acknowledgment delay to 600 milliseconds.

```
Router# configure terminal
Router(config)# router isis 1
Router(config-router)# lsp-fast-flooding
Router(config-isis-lsp-fast-flooding)# remote-psnp-delay 600
```

lsp-password

To configure the link-state packet (LSP) authentication password, use the **lsp-password** command in XR Config 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**]]

Syntax Description

hmac-md5	Specifies that the password uses HMAC-MD5 authentication.
text	Specifies that the password uses clear text password authentication.
clear	Specifies that the password be unencrypted.
encrypted	Specifies that the password be encrypted using a two-way algorithm.
password	Authentication password you assign.
keychain	(Optional) Specifies a keychain.
keychain-name	Name of the keychain.
level { **1** | **2** }	(Optional) Specifies the password for Level 1 or Level 2 independently.
send-only	(Optional) Adds passwords to LSP and sequence number protocol (SNP) data units when they are sent. Does not check for authentication in received LSPs or sequence number PDUs (SNPs).
snp send-only	(Optional) Adds passwords to SNP data units when they are sent. Does not check for authentication in received SNPs. This option is available when the **text** keyword is specified.
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

XR Config mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

When a **text** password is configured, it is exchanged as clear text. Therefore, the **lsp-password** command provides limited security.

When an **HMAC-MD5** password is configured, the password is never sent over the network and is instead used to calculate a cryptographic checksum to ensure the integrity of the exchanged data.

The recommended password configuration is that both incoming and outgoing SNPs be authenticated.

Note To disable SNP password checking, the **snp send-only** keywords must be specified in the **lsp-password** command.

To configure an additional password, use the **lsp-password accept** command.

Specify a key chain to enable key chain authentication between two IS-IS peers. Use the **keychain** *keychain-name* keyword and argument to implement hitless key rollover for authentication.

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
```

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 XR Config mode. To restore the default refresh interval, use the **no** form of this command.

lsp-refresh-interval *seconds* [**level** { **1** | **2** }]

Syntax Description

seconds	Refresh interval (in seconds). Range is 1 to 65535 seconds.
level { **1** | **2** }	(Optional) Specifies routing Level 1 or Level 2 independently.

Command Default

seconds : 900 seconds (15 minutes)

Both Level 1 and Level 2 are configured if no level is specified.

Command Modes

XR Config mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

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
```

maximum-paths (IS-IS)

To configure the maximum number of parallel routes that an IP routing protocol will install in the routing table, use the **maximum-paths** command in address family configuration mode. To remove the **maximum-paths** command from the configuration file and restore the system to its default condition with respect to the routing protocol, use the **no** form of this command.

maximum-paths *maximum*

Syntax Description

maximum	Maximum number of parallel routes that IS-IS can install in a routing table. Range is 1 to 64

Command Default

No default Command

Command Modes

Address family configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

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** }]

Syntax Description

maximum	Maximum number of redistributed prefixes advertised. Range is 1 to 28000.
level { **1** | **2** }	(Optional) Specifies maximum prefixes for Level 1 or Level 2.

Command Default

maximum: 10000

level : 1-2

Command Modes

Address family configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

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/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-redistributed-prefixes 5000 level 2
```

maximum-propagated-prefixes (IS-IS)

To configure the maximum number of propagated prefixes, use the **maximum-propagated-prefixes** command in mode.

maximum-propagated-prefixes *limit* [**warning-only** | **warning-percentage** *percentage*] [**level** *value*]

Syntax Description

maximum-propagated-prefixes *limit*	Configures the maximum number of propagated prefixes. The number of propagated prefixes ranges from 1 to 100000.
warning-only	Configures the warning-only behavior. The default warning-only threshold is 20,000 prefixes.
warning-percentage *percent*	Configures the warning threshold percentage. The default percentage is 70%. The warning threshold percentage ranges from 20 percentage and 99 percentage.
level *value*	Applies the limit to level 1 or level 2 propagation.

Command Default

Limit: 20,000

Behavior: warning-only

Warning percentage: 70% when a hard limit applies

Command Modes

IS-IS address-family configuration

Command History

Release	Modification
Release 26.3.1	This command was introduced.

Usage Guidelines

None

Task ID

Task ID	Operations
isis	read, write

Examples

This example shows how to configure the maximum propagated prefix limit for the level 2:

```
Router(config)# router isis isp
Router(config-isis)# address-family ipv4 unicast
Router(config-isis-af)# maximum-propagated-prefixes 8000 level 2
```

This example shows how to set an early warning threshold as a percentage.

```
Router(config)# router isis isp
Router(config-isis)# address-family ipv4 unicast
Router(config-isis-af)# maximum-propagated-prefixes 8000 warning-percentage 70 level 2
```

This example shows how to configure a warning-only threshold.

```
Router(config)# router isis isp
Router(config-isis)# address-family ipv4 unicast
Router(config-isis-af)# maximum-propagated-prefixes 20000 warning-only level 2
```

max-lsp-lifetime

To set the maximum time that link-state packets (LSPs) persist without being refreshed, use the **max-lsp-lifetime** command in XR Config mode. To restore the default time, use the **no** form of this command.

max-lsp-lifetime *seconds* [**level** { **1** | **2** }]

Syntax Description

seconds	Lifetime (in seconds) of the LSP. Range from 1 to 65535 seconds.
level { **1** | **2** }	(Optional) Specifies routing Level 1 or Level 2 independently.

Command Default

seconds : 1200 seconds (20 minutes)

Both Level 1 and Level 2 are configured if no level is specified.

Command Modes

XR Config mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

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
```

max-lsp-tx

To set the maximum number of Link State Packets (LSPs) that can be transmitted per second in an IS-IS routing process, use the **max-lsp-tx** command in IS-IS router lsp-fast-flooding configuration mode. To remove the configured maximum LSP transmission rate, use the **no** form of this command.

max-lsp-tx *lsps-per-sec*

Syntax Description

lsps-per-sec	The maximum number of LSPs that can be transmitted per second. The default is 1000 ms. The range is from 33 through 5000 ms.

Command Default

The command is disabled by default.

Command Modes

IS-IS router lsp-fast-flooding configuration

Command History

Release	Modification
Release 24.3.1	This command was introduced.

Usage Guidelines

None

Task ID

Task ID	Operation
isis	read, write

Example

This example shows how to set maximum LSP transmission rate to 600 LSPs per second.

```
Router# configure terminal
Router(config)# router isis 1
Router(config-isis)# lsp-fast-flooding
Router(config-isis-lsp-fast-flooding)# max-lsp-tx 600
```

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** }]

Syntax Description

default-metric	Metric assigned to the link and used to calculate the cost from each other router using the links in the network to other destinations. Range is 1 to 63 for narrow metric and 1 to 16777214 for wide metric. **Note** Setting the default metric under address family results in setting the same metric for all interfaces that is associated with the address family. Setting a metric value under an interface overrides the default metric
maximum	Specifies maximum wide metric. All routers exclude this link from their shortest path first (SPF).
level { **1** | **2** }	(Optional) Specifies the SPF calculation for Level 1 or Level 2 independently.

Command Default

default-metric : Default is 10.

Both Level 1 and Level 2 are configured if no level is specified.

Command Modes

Address family configuration

Interface address family configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

Specifying the **level** keyword resets the metric only for the specified level. We highly recommend that you configure metrics on all interfaces.

Set the default metric under address family to set the same metric for all interfaces that is associated with the address family. Set a metric value under an interface to override the default metric.

We highly recommend that you configure metrics on all interfaces.

Metrics of more than 63 cannot be used with narrow metric style.

Task ID

Task ID	Operations
isis	read, write

Examples

The following example shows how to configure HundredGigE 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 HundredGigE0/1/0/1
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
```

metric fallback anomaly

To configure the metric when a delay or loss anomaly occurs for an Intermediate System-to-Intermediate System (IS-IS) interface, use the **metric fallback anomaly** command in address family or interface address family configuration mode.

metric fallback anomaly { [**delay** | **loss**] [**te-metric**] { **increment** *penalty_value* | **multiplier** *penalty_multiplier* | **maximum** } } [**level** { **1** | **2** }]

Syntax Description

maximum	Specifies maximum wide metric. All routers exclude this link from their shortest path first (SPF).
delay	Increases the link metric only when a delay anomaly occurs.
loss	Increases the link metric only when a loss anomaly occurs.
te-metric	Modifies the TE metric.
increment *penalty_value*	Increments the link metric by the specified amount. The range is from 1 to 16777213.
multiplier *penalty_multiplier*	Multiplies the link metric by the specified amount. The range is from 2 to 100000.
level { **1** | **2** }	(Optional) Specifies the SPF calculation for Level 1 or Level 2 independently.

Command Default

Level 1 and Level 2 are configured if no level is specified.

Command Modes

Address family configuration

Interface address family configuration

Command History

Release	Modification
Release 25.1.1	The **loss** keyword was introduced.
Release 7.4.1	This command was introduced.

Usage Guidelines

Specifying the **level** keyword resets the metric only for the specified level. We highly recommend that you configure metrics on all interfaces.

Use the **fallback anomaly** options to increase the metric when you configure link delay anomaly detection in Segment Routing Performance Measurement (SR-PM).

When you configure link delay anomaly detection, SR-PM sends an anomaly bit (A-bit) when the upper bound of the anomaly threshold is exceeded. When an anomaly is received for an interface, IS-IS can increase the IGP metric or TE metric to make the link less attractive for traffic.

From Release 25.1.1, you can configure link loss anomaly detection so that SR-PM sends an anomaly bit (A-bit) when the upper bound of the anomaly threshold is exceeded. When an anomaly is received for an interface, IS-IS can increase the IGP metric or TE metric to make the link less attractive for traffic.

You can specify how to increase the metric in the following ways:

- As a factor by which the base metric is multiplied
- As an increment that is added to the base metric
- As the maximum value. All routers will exclude this link from their SPF.

Task ID

Task ID	Operations
isis	read, write

Examples

In the following example, the base IGP metric value is 10. This example shows how to increase the IGP metrics for an interface. The IGP metric will be increased by 50. When an anomalous bit notification comes for interface HundredGigE 0/1/0/0, the value advertised in the TLVs will be IGP metric 60.

```
Router(config)# router isis 1
Router(config-isis)# interface HundredGigE 0/1/0/0
Router(config-isis-if)# address-family ipv4 unicast
Router(config-isis-if-af)# metric fallback anomaly delay increment 50
```

From Release 25.1.1, you can configure link loss anomaly detection. This example displays how to increase the IGP metrics for an interface.

Consider that the base IGP metric value is 10. You can configure the IGP metric to increase by 500, so that if an anomalous bit notification occurs for interface HundredGigE 0/1/0/0 due to packet loss, the value advertised in the TLVs will be IGP metric 510.

```
Router(config)# router isis 1
Router(config-isis)# interface HundredGigE 0/1/0/0
Router(config-isis-if)# address-family ipv4 unicast
Router(config-isis-if-af)# metric fallback anomaly loss increment 500
```

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** }]

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 7.0.12	This command was introduced.

Usage Guidelines

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/RSP0RP0/CPU0:router(config-isis-af)# metric-style wide level 1
```

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

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 7.0.12	This command was introduced.

Usage Guidelines

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
```

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** }]

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 7.0.12	This command was introduced.

Usage Guidelines

MPLS VPN traffic forwarded using LDP labels can be dropped in the following instances:

- A new link is introduced in the network and IS-IS has converged before LDP establishes labels.
- An existing LDP session goes down while IS-IS adjacency is intact over the link.

In both instances, outbound LDP labels are not available for forwarding MPLS traffic. LDP IS-IS synchronization addresses the traffic drop. When the **mpls ldp sync** command is configured, IS-IS advertises the maximum possible link metric until LDP has converged over the link. The link is less preferred and least used in forwarding MPLS traffic. When LDP establishes the session and exchanges labels, IS-IS advertises the regular metric over the link.

Note IS-IS advertises the maximum metric –1 (16777214) if wide metrics are configured since the maximum wide metric is specifically used for link exclusion from the shortest path first algorithm (SPF) (RFC 3784). However, the maximum narrow metric is unaffected by this definition.

Task ID

Task ID	Operations
isis	read, write

Examples

The following example shows how to enable LDP IS-IS synchronization:

```
RP/0/RP0/CPU0:router(config)# router isis isp
RP/0/RP0/CPU0:router(config-isis)# interface HundredGigE 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
```

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** }

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

Usage Guidelines

Use the **mpls traffic-eng** command, which is part of the routing protocol tree, to flood link resource information (such as available bandwidth) for appropriately configured links in the link-state packet (LSP) of the router.

Task ID

Task ID	Operations
isis	read, write

Examples

The following example shows how to turn on MPLS traffic engineering for IS-IS level 1:

```
Router(config)# router isis isp
Router(config-isis)# address-family ipv4 unicast
Router(config-isis-af)# mpls  traffic-eng level-1
```

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* }

Syntax Description

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

Command Default

Global router identifier is used.

Command Modes

IPv4 address family configuration

Usage Guidelines

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:

```
Router(config)# router isis isp
Router(config-isis)# address-family ipv4 unicast
Router(config-isis-af)# mpls traffic-eng router-id Loopback0
```

oor-set-overload-bit disable

To prevent the setting of the overload bit during out-of-resource (OOR) conditions, use **oor-set-overload-bit disable** command in router configuration mode. To allow the setting of the overload bit during OOR conditions, use the **no** form of the command.

oor-set-overload-bit disable

no oor-set-overload-bit disable

Syntax Description

This command has no keywords or arguments.

Command Default

The overload bit is set.

Command Modes

Router configuration

Command History

Release	Modification
Release 24.2.1	This command was introduced.

Usage Guidelines

The setting of the overload bit is typically allowed only when a router experiences issues, such as memory shortages, which might result in an incomplete link-state database and inaccurate routing tables. When the overload bit is set in the link-state packets (LSPs) of an unreliable router, other routers can ignore this router in their shortest path first (SPF) calculations until it recovers. Consequently, no paths through the unreliable router are visible to other routers in the IS-IS area, though IP prefixes directly connected to this router remain reachable.

We recommend not disabling the setting of the overload-bit during OOR conditions. IS-IS will continue to forward traffic but will base decisions on an incomplete link-state database, potentially causing routing loops and other forwarding failures.

Task ID

Task ID	Operation
isis	read, write

Example

The following example shows how to disable the overload bit:

```
Router# router isis 1
Router# (config-isis)# oor-set-overload-bit disable
```

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*

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

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 95	Configures the routing level for an instance of the IS-IS routing process.
router isis, on page 135	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 XR Config mode. To restore the default setting, use the **no** form of this command.

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

XR Config mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

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
```

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

Command Default

Interface is active.

Command Modes

Interface configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

Task ID

Task ID	Operations
isis	read, write

Examples

The following example shows how to configure the router to suppress IS-IS packets on GigabitEthernet interface 0/1/0/1:

```
RP/0/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
```

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

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 7.0.12	This command was introduced.

Usage Guidelines

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 HundredGigE 0/6/0/0
RP/0/RP0/CPU0:router(config-isis-if)# point-to-point
```

protocol shutdown (IS-IS)

To disable an instance of the Intermediate System-to-Intermediate System (IS-IS) protocol, so that it cannot form an adjacency on any interface, use the **protocol shutdown** command in the isis configuration mode. To re-enable the IS-IS protocol, use the **no** form of this command.

protocol shutdown
no protocol shutdown

Command Default

None

Command Modes

isis configuration

Command History

Release	Modification
Release 24.4.1	This command was introduced.

Usage Guidelines

Use the **protocol shutdown** command to disable the IS-IS protocol for a specific routing instance without removing any existing IS-IS configuration parameters.

This command is similar to executing the following sample configuration:

```
router isis 1
 set-overload-bit
 exit
router isis 1
 interface HundredGigE 0/1/0/3
  shutdown
```

Task ID

Task ID	Operations
isis	read, write

Examples

This example shows how to disable the IS-IS.

```
Router# config
Router(config)# router isis 1
Router(config-isis)# protocol shutdown
```

psnp-interval

To configure the interval at which the Intermediate System to Intermediate System (IS-IS) protocol sends Partial Sequence Number PDUs (PSNPs) to acknowledge received Link State PDUs (LSPs), use the **psnp-interval** in router configuration mode.

psnp-interval *milliseconds*

Syntax Description

milliseconds	The interval, specified in milliseconds, at which PSNPs are sent to acknowledge received LSPs. The range is from 1 to 5000 milliseconds. The default value is 50 milliseconds.

Command Default

This command is disabled by default.

Command Modes

Router configuration

Command History

Release	Modification
Release 24.3.1	This command was introduced.

Usage Guidelines

None

Task ID

Task ID	Operation
isis	read, write

Example

This example shows how to configure the PSNP interval to 100 milliseconds for an IS-IS process tagged 1:

```
Router(config)# router isis 1
Router(config-isis)# psnp-interval 100
```

redistribute (IS-IS)

To redistribute routes from one routing protocol into Intermediate System-to-Intermediate System (IS-IS), use the **redistribute** command in address family configuration mode. To remove the **redistribute** command from the configuration file and restore the system to its default condition in which the software does not redistribute routes, use the **no** form of this command.

Border Gateway Protocol (BGP)
redistribute bgp *process-id* [**level-1** | **level-2** | **level-1-2**] [**metric** *metric-value*] [**metric-type** { **internal** | **external** | **rib-metric-as-external** | **rib-metric-as-internal** }] [**route-policy** *route-policy-name*]

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*]

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*] [**down-flag-clear**]

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*]

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*]

Static Routes
redistribute static [**level-1** | **level-2** | **level-1-2**] [**metric** *metric-value*] [**metric-type** { **1** { **internal** | **external** | **rib-metric-as-external** } | **2** **rib-metric-as-internal** }] [**route-policy** *route-policy-name*]

RIP
redistribute rip [**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*]

Subscriber Routes
redistribute subscriber [**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*]

Syntax Description

process-id	For the **bgp** keyword, an autonomous system number has the following ranges: • Range for 2-byte Autonomous system numbers (ASNs) is 1 to 65535. • Range for 4-byte Autonomous system numbers (ASNs) in asplain format is 1 to 4294967295. • Range for 4-byte Autonomous system numbers (ASNs) is asdot format is 1.0 to 65535.65535. For the **isis** keyword, an IS-IS instance identifier from which routes are to be redistributed. For the **ospf** keyword, an OSPF process name from which routes are to be redistributed. The value takes the form of a string. A decimal number can be entered, but it is stored internally as a string. For the **ospfv3** keyword, an OSPFv3 process name from which routes are to be redistributed. The value takes the form of a string. A decimal number can be entered, but it is stored internally as a string.
level-1	(Optional) Specifies that redistributed routes are advertised in the Level-1 LSP of the router.
level-1-2	(Optional) Specifies that redistributed routes are advertised in the Level-1-2 LSP of the router.
level-2	(Optional) Specifies that redistributed routes are advertised in the Level-2 LSP of the router.
metric *metric-value*	(Optional) Specifies the metric used for the redistributed route. Range is 0 to 16777215. The *metric-value* must be consistent with the IS-IS metric style of the area and topology into which the routes are being redistributed.
metric-type { **internal** | **external** } **metric-type** { **internal** | **external** | **rib-metric-as-external** | **rib-metric-as-internal** }	(Optional) Specifies the external link type associated with the route advertised into the ISIS routing domain. It can be one of two four values: • external • **internal** –Use the **internal** keyword to set IS-IS internal metric-type. • **external** –Use the **external** keyword to set IS-IS external metric-type • rib-metric-as-external–Use the **rib-metric-as-external** keyword to use RIB metric and set IS-IS external metric-type • rib-metric-as-internal–Use the **rib-metric-as-internal** keyword to use RIB metric and set IS-IS internal metric-type Any route with an internal metric (however large the metric is) is preferred over a route with external metric (however small the metric is). 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.

route-policy *route-policy-name*	(Optional) Specifies the identifier of a configured policy. A policy is used to filter the importation of routes from this source routing protocol to IS-IS.
match { **internal** | **external** [**1** | **2**] | **nsaa-external** [**1** | **2**]}	(Optional) Specifies the criteria by which OSPF routes are redistributed into other routing domains. It can be one or more of the following: • **internal** —Routes that are internal to a specific autonomous system (intra- and interarea OSPF routes). • **external** [**1 | 2**]—Routes that are external to the autonomous system, but are imported into OSPF as Type 1 or Type 2 external routes. • **nssa-external** [**1 | 2**]—Routes that are external to the autonomous system, but are imported into OSPF as Type 1 or Type 2 not-so-stubby area (NSSA) external routes. For the **external** and **nssa-external** options, if a type is not specified, then both Type 1 and Type 2 are assumed.
down-flag-clear	(Optional) Specifies that routes redistributed from another IS-IS instance should be advertised with the up/down bit set to zero. This is contrary to the behavior specified by RFC 5305 and RFC 7775 and can lead to route loops.

Command Default

Level 2 is configured if no level is specified.

metric-type: internal

match : If no match keyword is specified, all OSPF routes are redistributed.

Command Modes

Address family configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

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 command.

Only IPv4 OSPF addresses can be redistributed into IS-IS IPv4 address families and only IPv6 OSPFv3 prefixes can be distributed into IS-IS IPv6 address families.

Examples

In this example, IS-IS instance isp_A readvertises all of the routes of IS-IS instance isp_B in Level 2 LSP. Note that the **level-2** keyword affects which levels instance isp_A advertises the routes in and has no impact on which routes from instance isp_B are advertised. (Any Level 1 routes from IS-IS instance isp_B are included in the redistribution.

```
RP/0/RSP0RP0/CPU0:router(config)# router isis isp_A
RP/0/RSP0RP0/CPU0:router(config-isis)# net 49.1234.2222.2222.2222.00
RP/0/RSP0RP0/CPU0:router(config-isis)# address-family ipv4 unicast
RP/0/RSP0RP0/CPU0:router(config-isis-af)# redistribute isis isp_B level-2
!
RP/0/RSP0RP0/CPU0:router(config)# router isis isp_B
RP/0/RSP0RP0/CPU0:router(config-isis)# is-type level 1
RP/0/RSP0RP0/CPU0:router(config-isis)# net 49.4567.2222.2222.2222.00
RP/0/RSP0RP0/CPU0:router(config-isis)# address-family ipv4 unicast
```

remote-psnp-delay

To specify the maximum delay, in milliseconds, that the router should wait for a Partial Sequence Number Protocol (PSNP) acknowledgment from neighbors after sending a Link State Packet (LSP), use the **remote-psnp-delay** command in appropriate mode. To remove the configured maximum delay, use the **no** form of this command.

remote-psnp-delay *milliseconds*

Syntax Description

milliseconds	Maximum delay, in milliseconds, that the router should wait for a PSNP acknowledgment from neighbors after sending a LSP. The range is from 1 through 5000 ms.

Command Default

The command is disabled by default.

Command Modes

IS-IS router lsp-fast-flooding configuration

Interface configuration

Command History

Release	Modification
Release 24.3.1	This command was introduced.

Usage Guidelines

The order of precedence of the configuration using the **remote-psnp-delay** command is as follows:

1. The **remote-psnp-delay** command in interface configuration mode.
2. The value in the PSNP sub-TLV received from the neighbour, if any.
3. The **remote-psnp-delay** command in router configuration mode.

Task ID

Task ID	Operation
isis	read, write

Example

This example shows how to set the PSNP acknowledgment delay to 600 milliseconds on a device.

```
Router# configure terminal
Router(config)# router isis 1
Router(config-isis)# lsp-fast-flooding
Router(config-isis-lsp-fast-flooding)# remote-psnp-delay 600
Router(config-isis-lsp-fast-flooding)# commit
```

This example shows how to set the PSNP acknowledgment delay to 600 milliseconds on an interface.

```
Router# configure terminal
Router(config)# router isis 1
Router(config-isis)# interface GigabitEthernet0/0/0/0
Router(config-isis-if)# remote-psnp-delay 600
```

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 mode. To disable IS-IS routing, use the **no** form of this command.

router isis *instance-id*

Syntax Description

instance-id	Name of the routing process. Maximum number of characters is 40.

Command Default

An IS-IS routing protocol is not enabled.

Command Modes

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 122	Configures an IS-IS NET for the routing process.

router isis static neighbor

To enable the Intermediate System-to-Intermediate System (IS-IS) static neighbor and to configure its attributes, use the **router isis static-neighbor** command in XR Config mode. To disable IS-IS static neighbor, use the **no** form of this command.

router isis *instance-id* **static-neighbor**

Syntax Description

instance-id	Name of the routing process. Maximum number of characters is 40.

Command Default

None

Command Modes

XR Config mode

Command History

Release	Modification
25.2.1	This command was introduced.

Usage Guidelines

You can configure the following attributes by using the **static-neighbor** configuration mode.

system-id	ISIS System-id of the static neighbor.
remote ipv4/ipv6 address	Remote IPv4/IPv6 addresses.
affinity	Affinity link to advertise.
log advertisement changes	Enable logging static neighbor advertisement changes.

Task ID

Task ID	Operations
isis	read, write

Examples

The following example shows how to configure IS-IS static neighbor and its attributes for a router:

```
Router#config
Router(config)#router isis 1
Router(config-isis)#interface GigabitEthernet0/2/0/7
Router(config-isis-if)#static-neighbor
Router(config-isis-static-nbr)#system-id 0001.0002.0007
Router(config-isis-static-nbr)#remote ipv4 address 2.2.2.2
Router(config-isis-static-nbr)#affinity RED
Router(config-isis-static-nbr)#log-advertisement-changes
```

set default-isis-instance

When multiple IS-IS instances are configured on a router, by default, **show isis** commands display information from all IS-IS instances. To display information from only one specific IS-IS instance, use the **set default-isis-instance** *instance-name* command in the XR EXEC mode.

When the default IS-IS instance is set, all the show IS-IS commands during the terminal session display only information from that default IS-IS instance. To revert back to the default behaviour of displaying information from all IS-IS instances, use the **set default-isis-instance all** command.

set default-isis-instance { *instance-name* | **all** }

Syntax Description

instance-name	Displays information for the specified *instance-name* when running **show isis** commands.
all	Displays information from all configured IS-IS instances when running **show isis** commands.

Command Default

By default, all configured IS-IS instances are displayed when you run any **show isis** command.

Command Modes

XR EXEC mode

Usage Guidelines

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

Task ID

Task ID	Operations
isis	read, write

Examples

Example 1: This example shows how to set a specific IS-IS instance as the default instance.

```
RP/0/1/CPU0:ios#set default-isis-instance india
Tue Apr 23 11:47:04.353 CEST
%% Default isis instance is set to  india
```

Example 2: This example uses the **show isis summary** command to display information from the default instance set in Example 1.

```
RP/0/1/CPU0:ios#show isis summary
Tue Apr 23 11:47:10.802 CEST

                                (L1/L2)     (L1/L2)      (L1/L2)      (IPv4/IPv6)
Tag                    Level    Adj         Routers      LSPs         Routes
---                    -----    ---         -------      ----         ------
india                  L2       0/1         0/2          0/2          0/0
RP/0/1/CPU0:ios#
```

Example 3: This example shows how to set all configured IS-IS instances as default instances.

```
RP/0/1/CPU0:ios#set default-isis-instance all
Tue Apr 23 11:47:19.308 CEST
%% Default isis instance is set to  all
```

Example 4: This example uses the **show isis summary** command to display information from all the configured instances set in Example 3.

```
RP/0/1/CPU0:ios#show isis summary
Tue Apr 23 11:47:24.697 CEST

                                (L1/L2)     (L1/L2)      (L1/L2)      (IPv4/IPv6)
Tag                    Level    Adj         Routers      LSPs         Routes
---                    -----    ---         -------      ----         ------
eu                     L2       0/2         0/6          0/6          5/5
india                  L2       0/1         0/2          0/2          0/0
RP/0/1/CPU0:ios#
```

set-overload-bit

To configure the router to signal other routers not to use it as an intermediate hop in their shortest path first (SPF) calculations, use the **set-overload-bit** command in XR Config 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** }]

Syntax Description

on-startup	(Optional) Sets the overload bit only temporarily after reboot.
delay	(Optional) Time (in seconds) to advertise when the router is overloaded after reboot. Range is 5 to 86400 seconds (86400 seconds = 1 day).
wait-for-bgp	(Optional) Sets the overload bit on startup until the Border Gateway Protocol (BGP) signals converge or time out.
level { **1** | **2** }	(Optional) Specifies the overload bit for Level 1 or Level 2 independently.
advertise { **external** | **interlevel**	(Optional) Sets the overload bit set if the router advertises the following types of IP prefixes: • **external**—If overload-bit set advertises IP prefixes learned from other protocols • **interlevel**— If overload-bit set advertise IP prefixes learned from another ISI S level

Command Default

The overload bit is not set.

Both Level 1 and Level 2 are configured if no level is specified.

Command Modes

XR Config mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

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
```

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 mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

For each instance, the first line of output lists the IS-IS instance ID with the following lines identifying the IS-IS system ID, supported levels (level 1, level 2, or level-1-2), configured area addresses, active area addresses, status (enabled or not) and type (Cisco or IETF) of nonstop forwarding (NSF), and the mode in which the last IS-IS process startup occurred.

Next, the status of each configured address family (or just IPv4 unicast if none are configured) is summarized. For each level (level 1 or level 2), the metric style (narrow or wide) generated and accepted is listed along with the status of incremental shortest path first (iSPF) computation (enabled or not). Then redistributed protocols are listed, followed by the administrative distance applied to the redistributed routes.

Finally, the running state (active, passive, or disabled) and configuration state (active or disabled) of each IS-IS interface is listed.

Task ID

Task ID	Operations
isis	read

Examples

The following is sample output from the **show isis** command:

```
RP/0/RP0/CPU0:router# show isis
Wed Aug 20 23:54:55.043 PST DST

IS-IS Router: lab
  System Id: 0000.0000.0002
  IS Levels: level-2-only
  Manual area address(es):
    49.1122
```

```
Routing for area address(es):
  49.1122
Non-stop forwarding: Disabled
Most recent startup mode: Cold Restart
Topologies supported by IS-IS:
  IPv4 Unicast
    Level-2
      Metric style (generate/accept): Narrow/Narrow
      Metric: 10
      ISPF status: Disabled
    No protocols redistributed
    Distance: 115
Interfaces supported by IS-IS:
  Loopback0 is running passively (passive in configuration)
  POS0/1/0/2 is running actively (active in configuration)
  POS0/1/0/3 is running actively (active in configuration
```

This table describes the significant fields shown in the display.

Table 6: show isis Field Descriptions

Field	Description
IS-IS Router	IS-IS instance ID.
System Id	IS-IS system ID.
IS Levels	Supported levels for the instance.
Manual area address(es)	Domain and area.
Routing for area address(es):	Configured area addresses and active area addresses.
Non-stop forwarding	Status (enabled or not) and type (Cisco or IETF) of nonstop forwarding (NSF).
Most recent startup mode	The mode in which the last IS-IS process startup occurred.
Topologies supported by IS-IS	The summary of the status of each configured address family (or just IPv4 unicast if none are configured).
Redistributed protocols	List of redistributed protocols, followed by the administrative distance applied to the redistributed routes.
Metric style (generate/accept)	The status of each configured address family (or just IPv4 unicast if none are configured) is summarized. For each level (level 1 or level 2), the metric style (narrow or wide) generated and accepted is listed along with the status of incremental shortest path first (iSPF) computation (enabled or not).
Interfaces supported by IS-IS	The running state (active, passive, or disabled) and configuration state (active or disabled) of each IS-IS interface.

show isis adjacency

To display Intermediate System-to-Intermediate System (IS-IS) adjacencies, use the **show isis adjacency** command in XR EXEC mode.

show isis [**instance** *instance-id*] **adjacency** [**level** {**1** | **2**}] [*type interface-path-id*] [**detail**] [**systemid** *system-id*]

Syntax Description

instance *instance-id*	(Optional) Displays the IS-IS adjacencies for the specified IS-IS instance only. • The *instance-id* argument is the instance identifier (alphanumeric) defined by the **router isis** command.
level { **1** | **2** }	(Optional) Displays the IS-IS adjacencies for Level 1 or Level 2 independently.
type	Interface type. For more information, use the question mark (**?**) online help function.
interface-path-id	Physical interface or virtual interface. **Note** Use the **show interfaces** command to see a list of all interfaces currently configured on the router. For more information about the syntax for the router, use the question mark (**?**) online help function.
detail	(Optional) Displays neighbor IP addresses and active topologies.
systemid *system-id*	(Optional) Displays the information for the specified router only.

Command Default

No instance ID specified displays IS-IS adjacencies for all the IS-IS instances.

Both Level 1 and Level 2 are configured if no level is specified.

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

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 7: show isis adjacency Field Descriptions

Field	Description
Level-1	Level 1 adjacencies.
Level-2	Level 2 adjacencies.
System ID	Dynamic hostname of the system. The hostname is specified using the **hostname** command. If the dynamic hostname is not known or the **hostname dynamic disable** command has been executed, the 6-octet system ID is used.
Interface	Interface used to reach the neighbor.
SNPA	Data-link address (also known as the Subnetwork Point of Attachment [SNPA]) of the neighbor.
State	Adjacency state of the neighboring interface. Valid states are Down, Init, and Up.
Holdtime	Hold time of the neighbor.
Changed	Time the neighbor has been up (in hours:minutes:seconds).
NSF	Specifies whether the neighbor can adhere to the IETF-NSF restart mechanism.
BFD	Specifies the Bidirectional Forwarding Detection (BFD) status for the interface. Valid status are: • None—BFD is not configured. • Init—BFD session is not up. One reason is that other side is not yet enabled. • Up—BFD session has been established. • Down—BFD session holdtime expired.

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 *number* of entries. Range is 1 to 100.
first *number*	(Optional) Specifies that the output is restricted to the first *number* of entries. Range is 1 to 100.

Command Default

No default behavior or values

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

Task ID

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 8: 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

instance *instance-id*	(Optional) Displays the IS-IS checkpoint adjacencies for the specified IS-IS instance only. • The *instance-id* argument is the instance identifier (alphanumeric) defined by the **router isis** command.

Command Default

No instance ID specified displays IS-IS checkpoint adjacencies for all the IS-IS instances.

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

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
 Gi3/0/0/1           2    router-gsr8       Up     0001.0000.0008.04   80011fc4
```

This table describes the significant fields shown in the display.

Table 9: 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.

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

XR EXEC mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

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 10: 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.
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 *instance-id*	(Optional) Displays the IS-IS checkpoint LSPs for the specified instance only. • The *instance-id* argument is the instance identifier (alphanumeric) defined by the **router isis** command.

Command Default

No instance ID specified displays IS-IS checkpoint LSPs for all the IS-IS instances.

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

The checkpointed LSPs displayed by this command are used to restore the LSP database during a Cisco-proprietary nonstop forwarding (NSF) restart. The **show isis checkpoint lsp** command, with the **show isis database** command, may be used to verify the consistency of the two databases.

Task ID

Task ID	Operations
isis	read

Examples

The following is sample output from the **show isis checkpoint lsp** command:

```
RP/0/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 11: show isis checkpoint lsp Field Descriptions

Field	Description
Level	Routers with Level 1 or Level 2 adjacency configured.
LSPID	LSP identifier. The first six octets form the system ID of the router that originated the LSP. The next octet is the pseudonode ID. When this byte is 0 zero, the LSP describes links from the system. When it is nonzero, the LSP is a so-called nonpseudonode LSP. This is similar to a router link-state advertisement (LSA) in the Open Shortest Path First (OSPF) protocol. The LSP describes the state of the originating router. For each LAN, the designated router for that LAN creates and floods a pseudonode LSP, describing all systems attached to that LAN. The last octet is the LSP number. If there is more data than can fit in a single LSP, the LSP is divided into multiple LSP fragments. Each fragment has a different LSP number. An asterisk (*) indicates that the LSP was originated by the system on which this command is issued.
Chkpt ID	Unique ID issued to the checkpoint at its creation.

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 *instance-id*	(Optional) Displays the IS-IS LSP database for the specified instance only. • The *instance-id* argument is the instance identifier (alphanumeric) defined by the **router isis** command.
level { **1** | **2** }	(Optional) Displays the IS-IS LSP database for Level 1 or Level 2 independently.
update	(Optional) Displays contents of LSP database managed by update thread.
summary	(Optional) Displays the LSP ID number, sequence number, checksum, hold time, and bit information.
detail	(Optional) Displays the contents of each LSP.
verbose	(Optional) Displays the contents of each LSP.
* | *lsp-id*	(Optional) LSP protocol data units (PDUs) identifier. Displays the contents of a single LSP by its ID number or may contain an * as a wildcard character.

Command Default

No instance ID specified displays the IS-IS LSP database for all the IS-IS instances.

Both Level 1 and Level 2 is configured if no level is specified.

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 7.0.12	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 is sample output from the **show isis database** command with no keywords specified:

```
RP/0/RP0/CPU0:router# show isis database

IS-IS Area a1 (Level-1) Link State Database
LSPID                 LSP Seq Num  LSP Checksum  LSP Holdtime  ATT/P/OL
router-gsr6.00-00   * 0x00000016   0x62c8        896             0/0/0
router-gsr6.01-00   * 0x0000000f   0x56d9        902             0/0/0
router-gsr8.00-00     0x00000019   0x4b6d        1015            0/0/0
router-gsr9.00-00     0x00000016   0x33b7        957             0/0/0


 Total LSP count: 4 (L1: 4, L2 0, local L1: 2, local L2 0)
```

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.

```
RP/0/0/CPU0:ios#show isis database
Thu Dec 14 16:03:45.131 EST

IS-IS 100 (Level-1) Link State Database
LSPID                 LSP Seq Num   LSP Checksum  LSP Holdtime/Rcvd  ATT/P/OL
1111.1111.1111.00-00* 0x0073000e     0x0001                1002/*           1/0/1
1111.1111.1112.00-00  0x0073004a     0x0001                1195/59          1/0/1
1111.1111.1112.01-00  0x007300b8     0x0001                1178/59          0/0/1
1111.1111.1112.03-00  0x007300b6     0x0001                1179/59          0/0/1
1111.1111.1113.00-00  0x0073000d     0x0001                533/1200         1/0/1

 Total Level-1 LSP count: 5     Local Level-1 LSP count: 1
```

This table describes the significant fields shown in the display.

Table 12: show isis database Field Descriptions

Field	Description
LSPID	LSP identifier. The first six octets form the system ID of the router that originated the LSP. The next octet is the pseudonode ID. When this byte is 0, the LSP describes links from the system. When it is nonzero, the LSP is a so-called nonpseudonode LSP. This is similar to a router link-state advertisement (LSA) in the Open Shortest Path First (OSPF) protocol. The LSP describes the state of the originating router. For each LAN, the designated router for that LAN creates and floods a pseudonode LSP, describing all systems attached to that LAN. The last octet is the LSP number. If there is more data than can fit in a single LSP, the LSP is divided into multiple LSP fragments. Each fragment has a different LSP number. An asterisk (*) indicates that the LSP was originated by the system on which this command is issued.
LSP Seq Num	Sequence number for the LSP that allows other systems to determine if they have received the latest information from the source.
LSP Checksum	Checksum of the entire LSP packet.

Field	Description
LSP 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.
ATT/P/OL	ATT—Attach bit. This bit 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. P—P bit. Detects if the intermediate system is area partition repair capable. Cisco and other vendors do not support area partition repair. 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.

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 13: 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.

Field	Description
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.

The following is sample output from the **show isis database** command with the **detail verbose** keyword specified:

```
RP/0/RP0/CPU0:router# show isis instance isp database detail verbose

  IS-IS isp test (Level-1) Link State Database
      LSPID                 LSP Seq Num  LSP Checksum  LSP Holdtime  ATT/P/OL
      router-5router1.00-00       0x00000003 * 0x00000d44 0x8074460 0x7e2c 457 535
     0/0/0
        Area Address: 4900
    Area Address: 01
        NLPID:        0xcc
    NLPID:  Hostname:    router-50x8e
        IP Address:  172.1.1.5
        MetricMT: 0  IP 172.3.55.0/24Standard (IPv4 Unicast)
        Metric: 10         IP 172.6.1.0/24
    MT:           IPv6 Unicast                                       0/0/0
    MT:           IPv4 Multicast                  Metric: 10 IP 172.7.0./0/240
        Metric: 10         IS router-11.00
        Metric: 10         IS router-11.01
    MT:           IPv6 Multicast                router-11.00-00  * 0x0000000b   0x8074460
       1161            0/0/0
    Hostname:  Area Address: 49router1
        NLPID:       0xcc
        Hostname:    router-11
    IP Address:   192.168.0.145
        IP IPv6 Address:  172.1.11.11192:168::145     MetricRouter ID: 0        IP
172192.1168.1110.0/24145
        Metric: 10         IP 172IS-Extended router1.016.1.0/24
        Metric: 10         IP 172IS-Extended router2.007.0.0/24
        Metric: 10         IS routerIS-11Extended router2.0100
        Metric: 10         IS router-5.00
      router-11.01-00  * 0x00000001   0x80770ec        457             0/0/0
        Metric: 0          IS router-11.00
        Metric: 0          IS router-5.00
      Affinity: 0x00000000
      Interface IP Address: 10.3.11.145
      Neighbor IP Address: 10.3.11.143
      Physical BW: 155520 kbits/sec
       Total LSP count: 3 (L1: 3, L2 0, local L1: 2, local L2 0)
      Reservable Global pool BW: 0 kbits/sec
      Global Pool BW Unreserved:
      IS-IS isp (Level-2) Link State Database
      LSPID                 LSP Seq Num  LSP Checksum  LSP Holdtime  ATT/P/OL
      router-5.00-00     0x00000005   0x807997c        457             0/0/0
        [0]: 0        kbits/sec          [1]: 0        kbits/sec
        [2]: 0        kbits/sec          [3]: 0        kbits/sec
        [4]: 0        kbits/sec          [5]: 0        kbits/sec
        [6]: 0        kbits/sec          [7]: 0        kbits/sec
    MPLS SRLG: Area Address: 49router2.00
      Interface IP Address: 10.3.11.145
      Neighbor IP Address: 10.3.11.143
        NLPIDFlags:       0xcc0x1   HostnameSRLGs:    router-5IP Address[0]:  172.6.10,
[1.5]: 20
        Metric: 0 10  IP 172IP-Extended 10.3.5511.0/24
```

```
   Metric: 10         IP 172IP-Extended 192.1686.10.0145/2432
   Metric: 10         IS routerMT (IPv6 Unicast) IS-11Extended router1.0001
   Metric: 10         IP 172.1.0.0MT (IPv6 Unicast) IPv6 192:168::145/24128
   Metric: 10         IS routerMT (IPv4 Multicast) IS-11Extended router1.01
   Metric: 10         IP 172.8.111.0/24


 router-11.00-00  * 0x0000000d   0x807997c        1184            0/0/0
   Area Address: 49
   NLPID:       0xcc
   Hostname:    router-11
   IP Address:  172.28.111.111
   Metric: 0          IP 172.8.111.0/24
   Metric: 10         IP 172.6.1.0/24
   Metric: 10         IP 172MT (IPv4 Multicast) IP-Extended 192.7168.0./
   Metric: 10         IS router-11.01
   Metric: 10         IS router-5.00
   Metric: 10         IP 172.3.55.0MT (IPv6 Multicast) IPv6 192:168::145/24.01-00
0x0000013e 0x80770ec 0x3309 457 1159             0/0/0
   Metric: 0          IS routerIS-11Extended router1.00
   Metric: 0          IS routerIS-5Extended router2.00

  Total LSP count: 3 (L1: 0, L2 3, local L1: 0, local L2 2)
```

As the output shows, besides the information displayed with the **show isis database** command, the command with the **detail verbose** keyword displays the contents of each LSP.

Table 14: show isis instance isp database detail Field Descriptions

Field	Description
LSPID	LSP identifier. The first six octets form the system ID of the router that originated the LSP. The next octet is the pseudonode ID. When this byte is 0, the LSP describes links from the system. When it is nonzero, the LSP is a so-called nonpseudonode LSP. This is similar to a router link-state advertisement (LSA) in the Open Shortest Path First (OSPF) protocol. The LSP describes the state of the originating router. For each LAN, the designated router for that LAN creates and floods a pseudonode LSP, describing all systems attached to that LAN. The last octet is the LSP number. If there is more data than can fit in a single LSP, the LSP is divided into multiple LSP fragments. Each fragment has a different LSP number. An asterisk (*) indicates that the LSP was originated by the system on which this command is issued.
LSP Seq Num	Sequence number for the LSP that allows other systems to determine if they have received the latest information from the source.
LSP Checksum	Checksum of the entire LSP packet.
LSP Holdtime	Amount of 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.

Field	Description
ATT/P/OL	ATT—Attach bit. This bit 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. P—P bit. Detects if the intermediate system is area partition repair capable. Cisco and other vendors do not support area partition repair. 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.
Area Address	Reachable area addresses from the router. For Level 1 LSPs, these are the area addresses configured manually on the originating router. For Level 2 LSPs, these are all the area addresses for the area this route belongs to.
NLPID	Network Layer Protocol Identifier.
Hostname	Hostname of the node.
IP Address:	Address of the node.
Metric	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).

The following is additional sample output from the **show isis database detail** command. This is a Level 2 LSP. The area address 39.0001 is the address of the area in which the router resides.

```
RP/0/RSP0/CPU0:router# show isis database level 2 detail

IS-IS Level-2 Link State Database
LSPID                 LSP Seq Num  LSP Checksum  LSP Holdtime  ATT/P/OL
0000.0C00.1111.00-00* 0x00000006   0x4DB3        1194          0/0/0
  Area Address: 39.0001
  NLPID:       0x81 0xCC
  IP Address:  172.18.1.17
  Metric: 10   IS 0000.0C00.1111.09
  Metric: 10   IS 0000.0C00.1111.08
  Metric: 10   IP 172.17.4.0 255.255.255.0
  Metric: 10   IP 172.18.8.0 255.255.255.0
  Metric: 0    IP-External 10.0.0.0 255.0.0.0
```

The IP entries are the directly connected IP subnets the router is advertising (with associated metrics). The IP-External entry is a redistributed route.

Table 15: show isis database level 2 detail Field Descriptions

Field	Description
LSPID	LSP identifier. The first six octets form the system ID of the router that originated the LSP. The next octet is the pseudonode ID. When this byte is 0, the LSP describes links from the system. When it is nonzero, the LSP is a so-called nonpseudonode LSP. This is similar to a router link-state advertisement (LSA) in the Open Shortest Path First (OSPF) protocol. The LSP describes the state of the originating router. For each LAN, the designated router for that LAN creates and floods a pseudonode LSP, describing all systems attached to that LAN. The last octet is the LSP number. If there is more data than can fit in a single LSP, the LSP is divided into multiple LSP fragments. Each fragment has a different LSP number. An asterisk (*) indicates that the LSP was originated by the system on which this command is issued.
LSP Seq Num	Sequence number for the LSP that allows other systems to determine if they have received the latest information from the source.
LSP Checksum	Checksum of the entire LSP packet.
LSP 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.
ATT/P/OL	ATT—Attach bit. This bit 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. P—P bit. Detects if the intermediate system is area partition repair capable. Cisco and other vendors do not support area partition repair. 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.
Area Address	Reachable area addresses from the router. For Level 1 LSPs, these are the area addresses configured manually on the originating router. For Level 2 LSPs, these are all the area addresses for the area to which this route belongs.
NLPID	Network Layer Protocol Identifier.
Hostname	Hostname of the node.
IP Address:	IP address of the node.
Metric:	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).

Table 16: show isis database verbose Field Descriptions

Field	Description
LSPID	LSP identifier. The first six octets form the system ID of the router that originated the LSP. The next octet is the pseudonode ID. When this byte is zero, the LSP describes links from the system. When it is nonzero, the LSP is a so-called nonpseudonode LSP. This is similar to a router link-state advertisement (LSA) in the Open Shortest Path First (OSPF) protocol. The LSP describes the state of the originating router. For each LAN, the designated router for that LAN creates and floods a pseudonode LSP, describing all systems attached to that LAN. The last octet is the LSP number. If there is more data than can fit in a single LSP, the LSP is divided into multiple LSP fragments. Each fragment has a different LSP number. An asterisk (*) indicates that the LSP was originated by the system on which this command is issued.
LSP Seq Num	Sequence number for the LSP that allows other systems to determine if they have received the latest information from the source.
LSP Checksum	Checksum of the entire LSP packet.
LSP Holdtime	Time the LSP remains valid (in seconds). An LSP hold time of zero 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.
ATT/P/OL	ATT—Attach bit. This bit 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. P—P bit. Detects if the intermediate system is area partition repair capable. Cisco and other vendors do not support area partition repair. 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.
Area Address	Reachable area addresses from the router. For Level 1 LSPs, these are the area addresses configured manually on the originating router. For Level 2 LSPs, these are all the area addresses for the area to which this route belongs.
NLPID	Network Layer Protocol Identifier.
Hostname	Hostname of the node.
IP Address	IP address of the node.
Metric	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).

Field	Description
MPLS SRLG	MPLS SRLG TLV information per neighbor, identified by hostname or system ID.
Interface IP Address	Local interface IP address.
Neighbor IP Address	Remote interface IP address.
Flags	Flags carried in SRLG TLV. The Least Significant Bit (LSB) is set if the interface is numbered.
SRLGs	SRLG values.

This is the sample output from the **show isis database verbose** command. The output shows IPv4 adjacency segment ID (SID), prefix (node) SID, and Segment Routing Global Block (SRGB) values.

```
RP/0/RP0/CPU0:router show isis database verbose
Fri May  2 17:53:44.575 PDT

IS-IS DEFAULT (Level-1) Link State Database
LSPID                 LSP Seq Num  LSP Checksum  LSP Holdtime  ATT/P/OL
p1.00-00              0x00000080   0x4780        1044            1/0/0
  Area Address: 49.0001
  NLPID:        0xcc
  NLPID:        0x8e
  MT:           Standard (IPv4 Unicast)
  MT:           IPv6 Unicast                                     1/0/0
  Hostname:     p1
  IP Address:   172.16.255.101
  IPv6 Address: 2001:db8::ff:101
  Router Cap:   172.16.255.101, D:0, S:0
    Segment Routing: I:1 V:0, SRGB Base: 16000 Range: 7999
  Metric: 10         IS-Extended p2.00
    Interface IP Address: 172.16.2.4
    Neighbor IP Address: 172.16.2.5
    ADJ-SID: F:0 B:0 weight:0 Adjacency-sid:24002
  Metric: 10         IS-Extended pe1.00
    Interface IP Address: 172.16.1.1
    Neighbor IP Address: 172.16.1.0
    ADJ-SID: F:0 B:0 weight:0 Adjacency-sid:24003
  Metric: 10         IP-Extended 172.16.1.0/31
  Metric: 10         IP-Extended 172.16.2.2/31
  Metric: 10         IP-Extended 172.16.2.4/31
  Metric: 10         IP-Extended-Interarea 172.16.255.2/32
    Admin. Tag: 255
    Prefix-SID Index: 42, R:1 N:0 P:1
  Metric: 0          IP-Extended 172.16.255.101/32
    Prefix-SID Index: 141, R:0 N:0 P:0
  Metric: 10         MT (IPv6 Unicast) IS-Extended p2.00
  Metric: 10         MT (IPv6 Unicast) IS-Extended pe1.00
  Metric: 10         MT (IPv6 Unicast) IPv6 2001:db8::1:0/127
  Metric: 10         MT (IPv6 Unicast) IPv6 2001:db8::2:2/127
  Metric: 10         MT (IPv6 Unicast) IPv6 2001:db8::2:4/127
  Metric: 10         MT (IPv6 Unicast) IPv6-Interarea 2001:db8::ff:2/128
    Admin. Tag: 255
  Metric: 0          MT (IPv6 Unicast) IPv6 2001:db8::ff:101/128
```

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 *number* of entries. Range is 1 to 1000.
first *number*	(Optional) Specifies that the output be restricted to the first *number* of entries. Range is 1 to 1000.

Command Default

Both Level 1 and Level 2 are configured if no level is specified.

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

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
                                  New LSP                 Old LSP
WHEN     LSPID                Op  Seq Num    Holdtime OL  Seq Num    Holdtime OL
01:17:19 12b1.03-00           REP 0x00000003 1200      0  0x00000002 340       0
001:06:20 12b1.00-00          REP 0x000001d8 1200      0  0x000001d7 375       0
01:06:00 12b1.03-00           REP 0x00000004 1200      0  0x00000003 520       0
01:05:46 12a1.00-00           REP 0x000001fc 1200      0  0x000001fb 425       0
00:55:01 12b1.00-00           REP 0x000001d9 1200      0  0x000001d8 520       0
00:53:39 12b1.03-00           REP 0x00000005 1200      0  0x00000004 459       0
00:53:19 12a1.00-00           REP 0x000001fd 1200      0  0x000001fc 453       0
00:42:12 12b1.00-00           REP 0x000001da 1200      0  0x000001d9 431       0
00:39:56 12b1.03-00           REP 0x00000006 1200      0  0x00000005 376       0
00:38:54 12a1.00-00           REP 0x000001fe 1200      0  0x000001fd 334       0
00:29:10 12b1.00-00           REP 0x000001db 1200      0  0x000001da 418       0
00:27:22 12b1.03-00           REP 0x00000007 1200      0  0x00000006 446       0
00:25:10 12a1.00-00           REP 0x000001ff 1200      0  0x000001fe 375       0
```

```
00:17:04 12b1.00-00           REP 0x000001dc 1200      0  0x000001db 473
```

This table describes the significant fields shown in the display.

Table 17: show isis database-log Field Descriptions

Field	Description
WHEN	Elapsed time (in hh:mm:ss) since the event was logged.
LSPID	LSP identifier. The first six octets form the system ID of the router that originated the LSP. The next octet is the pseudonode ID. When this byte is 0 zero, the LSP describes links from the system. When it is nonzero, the LSP is a so-called nonpseudonode LSP. This is similar to a router link-state advertisement (LSA) in the Open Shortest Path First (OSPF) protocol. The LSP describes the state of the originating router. For each LAN, the designated router for that LAN creates and floods a pseudonode LSP, describing all systems attached to that LAN. The last octet is the LSP number. If there is more data than can fit in a single LSP, the LSP is divided into multiple LSP fragments. Each fragment has a different LSP number. An asterisk (*) indicates that the LSP was originated by the system on which this command is issued.
New LSP	New router or pseudonode appearing in the topology.
Old LSP	Old router or pseudonode leaving the topology.
Op	Operation on the database: inserted (INS) or replaced (REP).
Seq Num	Sequence number for the LSP that allows other systems to determine if they have received the latest information from the source.
Holdtime	Time the LSP remains valid (in seconds). An LSP hold time of 0 indicates that this LSP was purged and is being removed from the link-state database (LSDB) of all routers. The value indicates how long the purged LSP stays in the LSDB before being completely removed.
OL	Overload bit. Determines if the IS is congested. If the Overload bit is set, other routers do not use this system as a transit router when calculating routers. Only packets for destinations directly connected to the overloaded router are sent to this router.

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

A.B.C.D/length	Network to show per-prefix LFA information.
detail	Use to display tiebreaker information about the backup.
summary	Use to display the number of prefixes having protection per priority.
sr-only	Use to display SR-labeled prefixes only.

Command Default

None

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

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, HundredGigE 0/1/0/3, 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 frr IPv4 Unicast FRR summary

                                    Critical   High       Medium     Low        Total
                                    Priority   Priority   Priority      Priority

Prefixes reachable in L1
  All paths protected     0              0             2                8               10

  Some paths protected 0               0             1                3               4

  Unprotected                0              0             1                3
 4
  Protection coverage    0.00%     0.00%     75.00%     78.57%     77.78%
Prefixes reachable in L2
  All paths protected     0              0             0                0               0

  Some paths protected 0               0             1                0               1

  Unprotected                0              0             0                0
 0
  Protection coverage    0.00%     0.00%     100.00%   0.00%      100.00%
```

The following is sample output from **show isis fast-reroute sr-only** command that displays fast-reroute repair paths for prefixes associated with a segment routing prefix SID:

```
RP/0/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

instance *instance-id*	(Optional) Displays the IS-IS router name-to-system ID mapping table for the specified IS-IS instance only. The *instance-id* argument is the instance identifier (alphanumeric) defined by the **router isis** command.

Command Default

No instance ID specified displays the IS-IS router name-to-system ID mapping table for all the IS-IS instances.

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

The **show isis hostname** command does not display entries if the dynamic hostnames are disabled.

Task ID

Task ID	Operations
isis	read

Examples

The following is sample output from the **show isis hostname** command with the **instance** and *instance-id* values specified:

```
RP/0/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
```

This table describes the significant fields shown in the display.

Table 18: show isis instance isp hostname Field Descriptions

Field	Description
Level	IS-IS level of the router.

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.
Dynamic Hostname	Hostname of the router.
*	Local router.

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 mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.
Release 25.4.1	The output was modified to display the field *Bandwidth* and its value.

Usage Guidelines

No specific guidelines impact the use of this command.

Task ID

Task ID	Operations
isis	read

Examples

The following is sample output from the **show isis interface** command:

```
 Router# show isis interface
0/3/0/2
  tenGigE 0/3/0/2                 Enabled
   Adjacency Formation:       Enabled
   Prefix Advertisement:      Enabled
   Bandwidth:                 400000000

   BFD:                       Disabled
   BFD Min Interval:          150
```

```
BFD Multiplier:           3

Circuit Type:             level-2-only
Media Type:               P2P
Circuit Number:           0
Extended Circuit Number:  67111168
Next P2P IIH in:          4 s
LSP Rexmit Queue Size:    0

Level-2
  Adjacency Count:        1
  LSP Pacing Interval:    33 ms
  PSNP Entry Queue Size:  0

CLNS I/O
  Protocol State:         Up
  MTU:                    4469

IPv4 Unicast Topology:    Enabled
  Adjacency Formation:    Running
  Prefix Advertisement:   Running
  Metric (L1/L2):         80/80
  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 19: show isis interface Field Descriptions

Field	Description
Bandwidth: 400000000	Bandwidth of the interface, in kilobits per second (kbps). The system automatically detects this value, or you can manually configure it, in kilobits per second (kbps).
HundredGigE 0/6/0/0	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.

Field	Description
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).
LSP Rexmit Queue Size:	Number of LSPs pending retransmission on the interface.
Adjacency Count:	Number of adjacencies formed with a neighboring router that supports the same set of protocols.
PSNP Entry Queue Size:	Number of SNP entries pending inclusion in the next PSNP.
LAN ID:	ID of the LAN.
Priority (Local/DIS):	Priority of this interface or priority of the Designated Intermediate System.
Next LAN IIH in:	Time (in seconds) in which the next LAN hello message is sent.
LSP Pacing Interval:	Interval at which the link-state packet (LSP) transmission rate (and by implication the reception rate of other systems) is to be reduced.
Protocol State:	Running state of the protocol (up or down).
MTU:	Link maximum transmission unit (MTU).
SNPA:	Data-link address (also known as the Subnetwork Point of Attachment [SNPA]) of the neighbor.
All Level-n ISs:	Status of interface membership in Layer 2 multicast group. The status options are Yes or reason for not being a member of the multicast group.
IPv4 Unicast Topology:	Status of the topology, either enabled or disabled.
Adjacency Formation:	Status of adjacency formation. The status options are Running or a reason for not being ready to form adjacencies.
Prefix Advertisement:	Status of advertising prefixes, either enabled or disabled.
Metric (L1/L2):	IS-IS metric for the cost of the adjacency between the originating router and the advertised neighbor, or the metric of the cost to get from the advertising router to the advertised destination (which can be an IP address, an end system (ES), or a connectionless network service (CLNS) prefix).
MPLS LDP Sync (L1/L2)	Status of LDP IS-IS synchronization, either enabled or disabled. When enabled, the state of synchronization (Sync Status) is additionally displayed as either achieved or not achieved.
IPv4 Address Family:	Status of the address family, either enabled or disabled.
Protocol State:	State of the protocol.
Forwarding Address(es):	Addresses on this interface used by the neighbor for next-hop forwarding.

Field	Description
Global Prefix(es):	Prefixes for this interface included in the LSP.
LSP transmit timer expires in	LSP transmission expiration timer interval (in milliseconds).
LSP transmission is	State of LSP transmission. Valid states are: • idle • in progress • requested • requested and in progress

The following is sample output from the **show isis interface** command with the **brief** keyword:

```
RP/0/0/CPU0:router# show isis interface brief

    Interface      All     Adjs    Adj Topos  Adv Topos  CLNS   MTU    Prio
                   OK    L1   L2    Run/Cfg    Run/Cfg                L1   L2
-----------------  ---  ---------  ---------  ---------  ----  ----  --------
PO0/5/0/0          Yes    1    1      1/1        1/1      Up    4469    -    -
Gi0/6/0/0          Yes    1*   1*     1/1        1/1      Up    1497   64   64
```

This table describes the significant fields shown in the display.

Table 20: 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 *instance-id*	(Optional) Displays the LSP log information for the specified IS-IS instance only. • The *instance-id* argument is the instance identifier (alphanumeric) defined by the **router isis** command.
level { **1** | **2** }	(Optional) Displays the Intermediate System-to-Intermediate System (IS-IS) link-state database for Level 1 or Level 2 independently.
last *number*	(Optional) Specifies that the output be restricted to the last *number* of entries. Range is 1 to 20.
first *number*	(Optional) Specifies that the output be restricted to the first *number* of entries. Range is 1 to 20.

Command Default

No instance ID specified displays the LSP log information for all the IS-IS instances.

Both Level 1 and Level 2 are configured if no level is specified.

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

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 21: 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.

Field	Description
Triggers	A list of all reasons that triggered an LSP rebuild. The triggers are: • AREASET—area set changed • ATTACHFLAG—bit attached • CLEAR— **clear** command • CONFIG—configuration change • DELADJ—adjacency deleted • DIS—DIS changed • IFDOWN—interface down • IPADDRCHG—IP address change • IPDEFORIG—IP def-orig • IPDOWN—connected IP down • IFDOWN—interface down • IPEXT—external IP • IPIA—nterarea IP • IPUP—connected IP up • LSPDBOL—LSPDBOL bit • LSPREGEN—LSP regeneration • NEWADJ— new adjacency

show isis mesh-group

To display Intermediate System-to-Intermediate System (IS-IS) mesh group information, use the **show isis mesh-group** command in XR EXEC mode.

show isis [**instance** *instance-id*] **mesh-group**

Syntax Description

instance *instance-id*	(Optional) Displays the mesh group information for the specified IS-IS instance only. • The *instance-id* argument is the instance identifier (alphanumeric) defined by the **router isis** command.

Command Default

No instance ID specified displays the IS-IS mesh group information for all the IS-IS instances.

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

Task ID

Task ID	Operations
isis	read

Examples

The following is sample output from the **show isis mesh-group** command with the **instance** and *instance-id* values specified:

```
RP/0/RP0/CPU0:router# show isis instance isp mesh-group

 ISIS isp Mesh Groups

 Mesh group 6:
 HundredGigE 0/4/0/1
```

This table describes the significant fields shown in the display.

Table 22: 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 *instance-id*	(Optional) Displays the MPLS TE adjacency changes for the specified IS-IS instance only. • The *instance-id* argument is the instance identifier (alphanumeric) defined by the **router isis** command.
last *number*	(Optional) Specifies that the output is restricted to last *number* of entries. Range is 1 to 20.
first *number*	(Optional) Specifies that the output is restricted to first *number* of entries. Range is 1 to 20.

Command Default

No instance ID specified displays MPLS TE adjacency changes for all the IS-IS instances.

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

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 23: 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 179	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 *instance-id*	(Optional) Displays the latest flooded record from MPLS TE for the specified IS-IS instance only. • The *instance-id* argument is the instance identifier (alphanumeric) defined by the **router isis** command.

Command Default

No instance ID specified displays the latest flooded record from MPLS TE for all the IS-IS instances.

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

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 24: 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 177	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

instance *instance-id*	(Optional) Displays the MPLS TE tunnel information for the specified IS-IS instance only. • The *instance-id* argument is the instance identifier (alphanumeric) defined by the **router isis** command.

Command Default

No instance ID specified displays the MPLS TE tunnel information for all the IS-IS instances.

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

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

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 25: 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.

Field	Description
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 *instance-id*	(Optional) Displays the IS-IS neighbor information for the specified IS-IS instance only. • The *instance-id* argument is the instance identifier (alphanumeric) defined by the **router isis** command.
type	Interface type. For more information, use the question mark (**?**) online help function.
interface-path-id	Physical interface or virtual interface. **Note** Use the **show interfaces** command to see a list of all interfaces currently configured on the router. For more information about the syntax for the router, use the question mark (**?**) online help function.
summary	(Optional) Displays neighbor status count for each level.
detail	(Optional) Displays additional details.
systemid *system-id*	(Optional) Displays the information for the specified neighbor only.

Command Default

No instance ID specified displays neighbor information for all the IS-IS instances.

Both Level 1 and Level 2 are configured if no level is specified.

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

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/RSP0RP0/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
e333e LFA Interface: Gi0/1/0/0.1      0012.da6b.68a8 Up    8        L1   Capable 1
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 26: 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 27: 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.

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	*instance-id*	(Optional) Specifies the IS-IS 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 mode

Command History

Release	Modification
Release 26.3.1	The command output was modified to contain the details from the **maximum-propagated-prefixes** configuration.
Release 24.2.1	This command was modified. The fields to indicate the memory state of the IS-IS protocol were included.
Release 7.0.12	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

Task ID

Task ID	Operations
isis	read

Examples

This example shows the **show isis protocol** command output.

```
RP/0/RP0/CPU0:router# show isis protocol

IS-IS Router:                     eu
System Id:                    0000.0000.0005
Hostname:                     ios
IS Levels:                    level-1-2
Manual area address:          49.0002
Routing for area address:     49.0002
Multi-Instance Id:            0
Running on:                   RP/0/1/CPU0
Job Id:                       1009
PID:                          58446

Started:                      Tue 2026 Sep 08 14:42:26.442 CEST
  Respawn count:              3
  Null0 ready:                Tue 2026 Sep 08 14:42:26.680 CEST
  First adjacency up:         Tue 2026 Sep 08 14:42:28.029 CEST
  Protocol shutdown:          No
```

```
  Non-stop forwarding:           Disabled
  Most recent startup mode:      Cold Restart

Topologies supported by IS-IS
  IPv4 Unicast
    Distance:                    115
    Rib:                         connected
    Advertising local prefixes:  Passive Interfaces Only
    Redistributing:              no protocols
    Level-1
      Metric style:              Wide/Wide (generate/accept)
      Default Interface Metric:  10
      Advertising prefixes:      1
        Connected:               1
        Propagated:              0
          Warning when over:     20000
        Redistributed:           0
          Limit:                 10000
          Warning when over:     7000 (70%)
    Level-2
      Metric style:              Wide/Wide (generate/accept)
      Default Interface Metric:  10
      Advertising prefixes:      91
        Connected:               1
        Propagated:              90
          Warning when over:     20000
        Redistributed:           0
          Limit:                 10000
          Warning when over:     7000 (70%)
  IPv6 Unicast
    Distance:                    115
    Rib:                         connected
    Advertising local prefixes:  Passive Interfaces Only
    Redistributing:              no protocols
    Level-1
      Default Interface Metric:  10
      Advertising prefixes:      1
        Connected:               1
        Propagated:              0
          Warning when over:     10
        Redistributed:           0
          Limit:                 10000
          Warning when over:     7000 (70%)
    Level-2
      Default Interface Metric:  10
      Advertising prefixes:      18
        Connected:               1
        Propagated:              17
          Warning when over:     10 (exceeded)
        Redistributed:           0
          Limit:                 10000
          Warning when over:     7000 (70%)

Overload Bit
  Level-1:                       configured
    Current status:              clear
    Clear since:                 Tue 2026 Sep 08 14:47:26.791 CEST
    Configuration:               on startup for 300 seconds
      Interarea prefixes:        suppressed
      External prefixes:         suppressed
  Level-2:                       configured
    Current status:              clear
    Clear since:                 Tue 2026 Sep 08 14:47:26.789 CEST
    Configuration:               on startup for 300 seconds
```

```
      Interarea prefixes:         suppressed
      External prefixes:          suppressed
  Maximum Metrics:                not configured

  LSP MTU:                        1492
  LSP Full:                       level-1: No, level-2: No
  LSP fast flooding:              Disabled
    Max LSP transmit speed:       1000 LSPs/sec
    PSNP delay:                   500 ms
    PSNP interval:                50 ms

  SR-MPLS:
    SRLB allocated:               15000 - 15999
    SRGB (index 0) notified:      never
    SRGB (index 0):               allocated, 16000 - 23999

  MFI LSD registered:             Yes
    MFI LSD last registered:      Tue 2026 Sep 08 14:42:32.808 CEST
  TE connection status:           Up
    TE connection up:             Tue 2026 Sep 08 14:42:38.829 CEST
    TE connection flaps:          0
  XTC connection status:          Up
    XTC connection up:            Tue 2026 Sep 08 14:42:38.827 CEST
    XTC connection flaps:         0
  BFD connection status:          Up
    BFD connection up:            Tue 2026 Sep 08 14:42:27.600 CEST
    BFD connection flaps:         0
    BFD HA role change:           never
      BFD client HA role:         none
      BFD server HA role:         none

  System memory state:            Normal
    Last state change:            never
  IS-IS memory state:             Normal
    Current IS-IS memory usage:   22097440 bytes (1%)
    Limit (libmemdbg):            1610612400
    Last state change:            never
    Usage during process start:   733936 bytes (0%)
    Monitoring interval:          60 sec
  IS-IS memory usage per item:    items    bytes
    Nodes:                        7        2968
    LSPs:                         10       4459
    Links:                        28       34048
    Prefixes:                     385      373440
    Nets:                         139      108976
    Paths:                        334      32064

  Interfaces supported by IS-IS eu:
    Loopback0 is running passively (passive in configuration)
    GigabitEthernet0/0/0/35 is running actively (active in configuration)
    GigabitEthernet0/0/0/45 is running actively (active in configuration)
    GigabitEthernet0/0/0/56 is running actively (active in configuration)
    GigabitEthernet0/0/0/65 is running actively (active in configuration)
```

This example shows the output containing propagated prefix limits configured using the **maximum-propagated-prefixes** command.

```
Router# show isis protocol
IS-IS Router:                     eu
  System Id:                      0000.0000.0005
  Hostname:                       ios
  IS Levels:                      level-1-2
```

```
Manual area address:              49.0002
Routing for area address:         49.0002
Multi-Instance Id:                0
Running on:                       RP/0/1/CPU0
Job Id:                           1009
PID:                              58446

Started:                          Tue 2026 Sep 08 14:42:26.442 CEST
  Respawn count:                  3
  Null0 ready:                    Tue 2026 Sep 08 14:42:26.680 CEST
  First adjacency up:             Tue 2026 Sep 08 14:42:28.029 CEST
  Protocol shutdown:              No
  Non-stop forwarding:            Disabled
  Most recent startup mode:       Cold Restart

Topologies supported by IS-IS
  IPv4 Unicast
    Distance:                     115
    Rib:                          connected
    Advertising local prefixes:   Passive Interfaces Only
    Redistributing:               no protocols
    Level-1
      Metric style:               Wide/Wide (generate/accept)
      Default Interface Metric:   10
      Advertising prefixes:       26
        Connected:                1
        Propagated:               25
          Warning when over:      40
        Redistributed:            0
          Limit:                  10000
          Warning when over:      7000 (70%)
    Level-2
      Metric style:               Wide/Wide (generate/accept)
      Default Interface Metric:   10
      Advertising prefixes:       61
        Connected:                1
        Propagated:               60
          Limit:                  60 (exceeded)
          Warning when over:      30 (50%)
        Redistributed:            0
          Limit:                  10000
          Warning when over:      7000 (70%)
  IPv6 Unicast
    Distance:                     115
    Rib:                          connected
    Advertising local prefixes:   Passive Interfaces Only
    Redistributing:               no protocols
    Level-1
      Default Interface Metric:   10
      Advertising prefixes:       1
        Connected:                1
        Propagated:               0
          Warning when over:      10
        Redistributed:            0
          Limit:                  10000
          Warning when over:      7000 (70%)
    Level-2
      Default Interface Metric:   10
      Advertising prefixes:       18
        Connected:                1
        Propagated:               17
          Warning when over:      10 (exceeded)
        Redistributed:            0
          Limit:                  10000
```

```
          Warning when over:     7000 (70%)
Overload Bit:                    not configured
Maximum Metrics:                 not configured

LSP MTU:                         1492
LSP Full:                        level-1: No, level-2: No
LSP fast flooding:               Disabled
  Max LSP transmit speed:        1000 LSPs/sec
  PSNP delay:                    500 ms
  PSNP interval:                 50 ms

SR-MPLS:
  SRLB allocated:                15000 - 15999
  SRGB (index 0) notified:       never
  SRGB (index 0):                allocated, 16000 - 23999

MFI LSD registered:              Yes
  MFI LSD last registered:       Tue 2026 Sep 08 14:42:32.808 CEST
TE connection status:            Up
  TE connection up:              Tue 2026 Sep 08 14:42:38.829 CEST
  TE connection flaps:           0
XTC connection status:           Up
  XTC connection up:             Tue 2026 Sep 08 14:42:38.827 CEST
  XTC connection flaps:          0
BFD connection status:           Up
  BFD connection up:             Tue 2026 Sep 08 14:42:27.600 CEST
  BFD connection flaps:          0
  BFD HA role change:            never
    BFD client HA role:          none
    BFD server HA role:          none

System memory state:             Normal
  Last state change:             never
IS-IS memory state:              Normal
  Current IS-IS memory usage:    22073984 bytes (1%)
  Limit (libmemdbg):             1610612400
  Last state change:             never
  Usage during process start:    733936 bytes (0%)
  Monitoring interval:           60 sec
IS-IS memory usage per item:     items    bytes
  Nodes:                         7        2968
  LSPs:                          10       4454
  Links:                         28       34048
  Prefixes:                      380      370912
  Nets:                          139      108976
  Paths:                         334      32064

Interfaces supported by IS-IS eu:
  Loopback0 is running passively (passive in configuration)
  GigabitEthernet0/0/0/35 is running actively (active in configuration)
  GigabitEthernet0/0/0/45 is running actively (active in configuration)
  GigabitEthernet0/0/0/56 is running actively (active in configuration)
  GigabitEthernet0/0/0/65 is running actively (active in configuration)
```

show isis purge-log

To display the IS-IS LSP purge log to help troubleshoot purge events, use the **show isis purge-log** command in XR EXEC mode.

```
show isis purge-log [ last number | first number ][ level { 1 | 2 } ]
```

Table 28: Syntax Description

show isis purge-log	Executes the command without any additional parameters, displaying the default IS-IS purge log output.
first *number*	(Optional) Displays the first specified *number* of entries from the IS-IS purge log. The *number* can be from 1 to 50,000.
last *number*	(Optional) Displays the last specified *number* of entries from the IS-IS purge log. The *number* can be from 1 to 50,000.
level { **1** | **2** }	(Optional) Filters the display to show purge log entries for a specific IS-IS level, for example, level 1 or level 2.

Command Default

Displays the complete IS-IS purge log output without any filters.

Command Modes

XR EXEC mode

Table 29: Command History

Release	Modification
Release 25.4.1	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

Task ID	Operations
isis	read

This is a sample output of the **show isis purge-log** command.

```
Router#show isis purge-log

  IS-IS 1  Level 2  LSP Purge Log                                  Capacity: 500, Size:
4

When          LSP ID                 Received from    On            POI originator  POI
neighbor
--- Mon Jun 16 2025 ---
12:56:04.295  r104.01-00             r100             Gi0/0/0/1     r100            r103
12:56:05.297  r103.01-00             r100             Gi0/0/0/1     r100            r103
12:56:05.689  r104.01-00             r102             Gi0/0/0/2     r102            r103
12:56:07.297  r103.01-00             r102             Gi0/0/0/2     r102            r103
```

This table describes the significant fields shown in the display.

Table 30: Field Descriptions

Field	Description
When	Time when the LSP purge was logged.
LSP ID	Identifier of the LSP including originating router and LSP number.
Received from	Router ID of the neighbor that sent the purged LSP.
On	Interface on which the LSP was received.
POI originator	Router ID of the original LSP originator.
POI neighbor	Router ID of the neighbor related to the purge event.

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** | **multicast** [**topology** {**all***topo-name*}] | **safi-all**] **route** [*ip-address mask* | *ip-address*/*length* [**longer-prefixes**]] [**summary**] **[multicast-intact]** [**backup**] [**detail**] [**sr-only**]

Syntax Description

instance *instance-id*	(Optional) Displays the IP reachability information for the specified IS-IS instance only. The *instance-id* argument is the instance identifier (alphanumeric) defined by the **router isis** command.
ipv4	(Optional) Specifies IP Version 4 address prefixes.
ipv6	(Optional) Specifies IP Version 6 address prefixes.
afi-all	(Optional) Specifies all address prefixes.
unicast	(Optional) Specifies unicast address prefixes.
multicast	(Optional) Specifies multicast address prefixes.
topology	(Optional) Specifies IS-IS paths to intermediate systems.
all	(Optional) Specifies all topologies.
topology *topo-name*	(Optional) Specifies topology table information and name of the topology table.
safi-all	(Optional) Specifies all secondary address prefixes.
ip-address	(Optional) Network IP address about which routing information should be displayed.
mask	(Optional) Network mask specified in either of two ways: • Network mask can be a four-part, dotted decimal address. For example, 255.0.0.0 indicates that each bit equal to 1 means the corresponding address bit is a network address. • Network mask can be indicated as a slash (/) and number. For example, /8 indicates that the first 8 bits of the mask are ones, and the corresponding bits of the address are the network address.
/ *length*	(Optional) Length of the IP prefix. A decimal value that indicates how many of the high-order contiguous bits of the address compose the prefix (the network portion of the address). A slash must precede the decimal value. Range is 0 to 32 for IPv4 prefix and 0 to 128 for IPv6 prefix.
longer-prefixes	(Optional) Displays route and more-specific routes.
summary	(Optional) Displays topology summary information.

multicast-intact	(Optional) Displays multicast intact information for this entry.
systemid	(Optional) Displays multicast information by system ID.
backup	(Optional) Displays 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 mode

Command History

Release	Modification
Release 26.3.1	The command output was modified for a route that was not propagated or advertised for over-limit.
Release 7.0.12	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

Task ID

Task ID	Operations
isis	read

Examples

The following is sample output from the **show isis route** command:

```
RP/0/RSP0RP0/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, HundredGigE 0/3/0/0, isp3
L2 10.76.240.10/32 [296/115]
via 10.76.249.2, HundredGigE 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, HundredGigE 0/3/0/0
  L2 10.101.10.0/24 [296/115]
  via 10.76.249.2, HundredGigE 0/3/0/0, isp3
```

This table describes the significant fields shown in the display.

Table 31: show isis route ipv4 unicast Field Descriptions

Field	Description
C172.18.0.0/24	Connected route for HundredGigE interface 0/5/0/0.
C 172.19.1.0/24	Connected route for HundredGigE 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 **detail** keyword that shows prefix segment ID (SID) and Segment Routing Global Block (SRGB) values:

```
Sun May  4 13:05:11.073 PDT

L2 172.16.255.2/32 [10/115] medium priority
     via 172.16.2.2, HundredGigE 0/0/0/1, pe2 tag 255, SRGB Base: 16000, Weight: 0
     src pe2.00-00, 172.16.255.2, tag 255, prefix-SID index 42, R:0 N:0 P:0
   L1 adv [10] native, propagated, interarea, tag 255, prefix-SID index 42, R:0
          N:0 P: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/RP0/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
       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
```

This example shows the output for a route that is not propagated or advertised due to an exceeded limit.

```
Router# show isis route 2.4.1.197/32 detail
Mon Aug 24 12:54:55.932 CEST

L1 2.4.1.197/32 [110/115] Label: None, medium priority
   Installed Aug 24 12:48:46.944 for 00:06:10
     via 1.1.45.4, GigabitEthernet0/0/0/45, brat, SRGB Base: 16000, Weight: 0
     src brat.00-02, 1.1.100.4
   L2 adv [110] native, propagated, over limit: not advertised
Router#
```

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** | **multicast** [**topology** {**all***topo-name*}] | **safi-all**]] **spf-log** [**level** {**1** | **2**}] [**ispf** | **fspf** | **prc** | **nhc**] [**detail** | **verbose** | **plfrr** | **ppfrr**] [**last** *number* | **first** *number*]

Syntax Description

instance *instance-id*	(Optional) Displays the IS-IS SPF log for the specified IS-IS instance only.
ipv4	(Optional) Specifies IP Version 4 address prefixes.
ipv6	(Optional) Specifies IP Version 6 address prefixes.
afi-all	(Optional) Specifies all address prefixes.
unicast	(Optional) Specifies unicast address prefixes.
multicast	(Optional) Specifies multicast address prefixes.
topology all | *topo-name*	(Optional) Specifies topology table information for all topologies or for the specified topology table (*top-name*).
safi-all	(Optional) Specifies all secondary address prefixes.
level { **1** | **2** }	(Optional) Displays the IS-IS SPF log for Level 1 or Level 2 independently.
ispf	(Optional) Specifies incremental SPF entries only.
fspf	(Optional) Specifies full SPF entries only.
prc	(Optional) Specifies partial route calculations only.
nhc	(Optional) Specifies next-hop route calculations only.
plfrr	(Optional) Specifies per link fast-reroute calculations only.
ppfrr	(Optional) Specifies per prefix fast-reroute calculations only.
detail	(Optional) Specifies detailed output. Includes a breakdown of the time taken to perform the calculation and changes resulting from the calculation.
verbose	(Optional) Specifies verbose output.
last *number*	(Optional) Specifies that the output is restricted to the last *number* of entries. Range is 1 to 210.
first *number*	(Optional) Specifies that the output is restricted to the first *number* of entries. Range is 1 to 210.

Command Default

No instance ID specified displays IS-IS adjacencies for all the IS-IS instances.

Both Level 1 and Level 2 are configured if no level is specified.

Displays all types of route calculation (not just fspf, ispf and prc).

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

Task ID

Task ID	Operations
isis	read

Examples

The following is sample output from the **show isis spf-log** command:

```
RP/0/RP0/CPU0: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 32: show isis spf-log ipv4 unicast Field Descriptions

Field	Description
Level	IS-IS level of the router.
Timestamp	Time when the SPF calculation started.
Duration	Number of milliseconds taken to complete this SPF run. Elapsed time is wall clock time, not CPU time.

Field	Description
Nodes	Number of routers and pseudonodes (LANs) that make up the topology calculated in this SPF run.
Trig Count	Number of events that triggered this SPF run. When there is a topology change, often multiple link-state packets (LSPs) are received in a short time. Depending on the configuration of the **spf-interval** command, a router may wait for a fixed period of time before running a router calculation. This count denotes the number of triggering events that occurred while the router was waiting to run the calculation. For a full description of the triggering events, see *List of Triggers*.
First Trigger LSP	LSP ID stored by the router whenever a full SPF calculation is triggered by the arrival of a new LSP. The LSP ID can suggest the source of routing instability in an area. If multiple LSPs are causing an SPF run, only the LSP ID of the first received LSP is remembered.
Triggers	List of all reasons that triggered a full SPF calculation. For a list of possible triggers, see *List of Triggers*.

This table lists triggers of a full SPF calculation.

Table 33: List of Triggers

Trigger	Description
PERIODIC	Runs a full SPF calculation very 15 minutes.
NEWLEVEL	Configured new level (using is-type) on this router.
RTCLEARED	Cleared IS-IS topology on the router.
MAXPATHCHANGE	Changed IP maximum parallel path.
NEWMETRIC	Changed link metric.
ATTACHFLAG	Changed Level 2 Attach bit.
ADMINDIST	Configured another administrative distance for the IS-IS instance on this router.
NEWADJ	Created a new adjacency to another router.
DELADJ	Deleted adjacency.
BACKUP	Installed backup route.
SEEDISPF	Seed incremental SPF.
NEXTHOP	Changed IP next-hop address.
NEWLSP0	New LSP 0 appeared in the topology.
LSPEXPIRED	Some LSP in the link-state database (LSDB) has expired.
LSPHEADER	Changed important LSP header fields.

Trigger	Description
TLVCODE	Type, length, and value (TLV) objects code mismatch, indicating that different TLV objects are included in the newest version of an LSP.
LINKTV	Changed Link TLV content.
PREFIXTLV	Changed Prefix TLV content.
AREAADDRTLV	Changed Area address TLV content.
IP ADDRTLV	Changed IP address TLV content.
TUNNEL	Changed RRR tunnel.

The following is sample output from the **show isis spf-log** command with the **first** keyword specified:

```
RP/0/RP0/CPU0:router# show isis spf-log first 2

    IISIS isp Level 1 IPv4 Unicast Route Calculation Log
                 Time  Total Trig
Timestamp   Type (ms)  Nodes Count First Trigger LSP   Triggers
  Mon Aug 16 2004
19:25:35.140  FSPF  1    1     1             12a5.00-00 NEWLSP0
19:25:35.646  FSPF  1    1     1                         NEWADJ


   IISIS isp Level 2  IPv4 Unicast Route Calculation Log
                 Time  Total Trig
Timestamp   Type (ms)  Nodes Count First Trigger LSP   Triggers
  Mon Aug 16 2004
19:25:35.139  FSPF  1    1     1             12a5.00-00 NEWLSP0
19:25:35.347  FSPF  1    1     2             12a5.00-00 NEWSADJ TLVCODE
```

This table describes the significant fields shown in the display.

Table 34: show isis spf-log first Field Descriptions

Field	Description
Level	IS-IS level of the router.
Timestamp	Time at which the SPF calculation started.
Type	Type of route calculation. The possible types are incremental SPF (iSPF), full SPF (FSPF), or partial route calculation (PRC).
Time (ms)	Number of milliseconds taken to complete this SPF run. Elapsed time is wall clock time, not CPU time.
Nodes	Number of routers and pseudonodes (LANs) that make up the topology calculated in this SPF run.

Field	Description
Trig Count	Number of events that triggered this SPF run. When there is a topology change, often multiple link-state packets (LSPs) are received in a short time. Depending on the configuration of the **spf-interval** command, a router may wait for a fixed period of time before running a router calculation. This count denotes the number of triggering events that occurred while the router was waiting to run the calculation. For a full description of the triggering events, see *List of Triggers*.
First Trigger LSP	LSP ID stored by the router whenever a full SPF calculation is triggered by the arrival of a new LSP. The LSP ID can suggest the source of routing instability in an area. If multiple LSPs are causing an SPF run, only the LSP ID of the first received LSP is remembered.
Triggers	List of all reasons that triggered a full SPF calculation. For a list of possible triggers, see *List of Triggers*.

The following is sample output from the **show isis spf-log** command with the **detail** keyword specified:

```
RP/0/RP0/CPU0: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 35: show isis spf-log detail Field Descriptions

Field	Description
Level	IS-IS level of the router.
Timestamp	Time at which the SPF calculation started.
Type	Type of route calculation. The possible types are incremental SPF (iSPF), full SPF (FSPF), or partial route calculation (PRC).
Time (ms)	Number of milliseconds taken to complete this SPF run. Elapsed time is wall clock time, not CPU time.
Nodes	Number of routers and pseudonodes (LANs) that make up the topology calculated in this SPF run.
Trig Count	Number of events that triggered this SPF run. When there is a topology change, often multiple link-state packets (LSPs) are received in a short time. Depending on the configuration of the **spf-interval** command, a router may wait for a fixed period of time before running a router calculation. This count denotes the number of triggering events that occurred while the router was waiting to run the calculation. For a full description of the triggering events, see *List of Triggers*.
First Trigger LSP	LSP ID stored by the router whenever a full SPF calculation is triggered by the arrival of a new LSP. The LSP ID can suggest the source of routing instability in an area. If multiple LSPs are causing an SPF run, only the LSP ID of the first received LSP is remembered.
Triggers	List of all reasons that triggered a full SPF calculation. For a list of possible triggers, see *List of Triggers*.
Delay	Two different delays exist: **1.** The delay between the time when the route calculation was first triggered and the time when it was run. **2.** The delay between the end of the last route calculation and the start of this one. This is used to verify that the SPF-interval timers are working correctly, and is only reported for calculations after the first delay.
CPU Time	Two different CPU times exist: **1.** CPU time (in milliseconds) taken to calculate the shortest path tree (SPT). **2.** CPU time (in milliseconds) taken to perform the prefix updates.
Real Time	Two different real times exist: **1.** Real time (in milliseconds) taken to calculate the shortest path tree (SPT). **2.** Real time (in milliseconds) taken to perform the prefix updates.

Field	Description
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.
Total	Total number of nodes or prefixes at various priorities.

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 *instance-id*	(Optional) Displays the IS-IS traffic statistics for the specified IS-IS instance only. • The *instance-id* argument is the instance identifier (alphanumeric) defined by the **router isis** command.
type	Interface type. For more information, use the question mark (**?**) online help function.
interface-path-id	Physical interface or virtual interface. **Note** Use the **show interfaces** command to see a list of all interfaces currently configured on the router. For more information about the syntax for the router, use the question mark (**?**) online help function.

Command Default

No instance ID specified displays IS-IS traffic statistics for all the IS-IS instances.

IS-IS traffic statistics are displayed for all interfaces.

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

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:

```
RP/0/RP0/CPU0: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
        Level-1::LSPs sourced (new/refresh)SPF calculations          : 3376/2754520   ISPF
 calculations        IPv4 Unicast: 0SPF calculations Next Hop Calculations            :
5200ISPF calculations          Partial Route Calculations : 0
          NextIPFRR R-hop Calculations   SPF calculations    : 0
          Partial Route Calculations IPFRR Parallel calculations: 0
        IPv6 Unicast
          SPF calculations           : 527
          ISPF calculations          : 0
          Next Hop Calculations      : 13
          Partial Route Calculations : 1
          Level-2:
        LSPs sourced (new/refresh): 4255/3332
        IPv4 Unicast
          SPF calculations           : 432
          ISPF calculations          : 0
          Next Hop Calculations      : 8
          LSPs sourced (new/refresh)Partial Route Calculations: 4255/33320LSPs sourced
(new/refresh)IPFRR R-SPF calculations   : 4255/33320
          IPFRR Parallel calculations: 0
        IPv4 IPv6 Unicast
          SPF calculations           : 432444
          ISPF calculations          : 0
          Next-hop Next Hop Calculations      : 882
          Partial Route Calculations : 01      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 36: 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.
Fast CSNP cache updates:	Number of times the CSNP cache has been updated since the last clearing of statistics. The cache is updated on LSP addition or removal from the database.
LSP checksum errors received:	Number of internal checksum errors received in LSPs.
IIH (LSP/SNP) dropped:	Number of hello, LSP, and SNP messages dropped.
IIH (UPD) Max Queue size:	Maximum number of queued packets.
Average transmit times and rate:	Average time taken to transmit the pdu type across all interfaces and the corresponding rate at which the pdu type is being transmitted.
Average process times and rate:	Average time taken to process an incoming pdu type across all interfaces and the corresponding rate at which the pdu type is being received.
LSPs sourced (new/refresh):	Number of LSPs this IS-IS instance has created or refreshed. To find more details on these LSPs, use the **show isis lsp-log** command.
SPF calculations:	Number of shortest path first (SPF) calculations. SPF calculations are performed only when the topology changes. They are not performed when external routes change. The interval at which SPF calculations are performed is configured using the **spf-interval** command.
iSPF calculations:	Number of incremental shortest path first (iSPF) calculations. iSPF calculations are performed only when ISPF has been configured in the isis address family configuration submode.
Partial Route Calculations:	Number of partial route calculations (PRCs). PRCs are processor intensive. Therefore, it may be useful to limit their number, especially how often a PRC is done, especially on slower networking devices. Increasing the PRC interval reduces the processor load on the router, but might slow the rate of convergence. The interval at which PRC calculations are performed is configured using the **spf-interval** command.
Level-(1/2) (LSPs/CSNPs/PSNPs/Hellos) (sent/rcvd):	Number of LSPs, Complete Sequence Number Packets (CSNPs), Partial Sequence Number Packets (PSNPs), and hello packets sent or received on this interface.

Field	Description
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.

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** | **multicast** [**topology** {**all** | **topo-name**}] | **safi-all**]] | **summary** | **level** {**1** | **2**} **[multicast-intact]** [**systemid system-id**] [**detail**]

Syntax Description

instance *instance-id*	(Optional) Displays the IS-IS topology for the specified IS-IS instance only. • The *instance-id* argument is the instance identifier (alphanumeric) defined by the **router isis** command.
ipv4	(Optional) Specifies IP Version 4 address prefixes.
ipv6	(Optional) Specifies IP Version 6 address prefixes.
afi-all	(Optional) Specifies all address prefixes.
unicast	(Optional) Specifies unicast address prefixes.
multicast	(Optional) Specifies multicast address prefixes.
topology *topo-name*	(Optional) Specifies topology table information and name of the topology table.
safi-all	(Optional) Specifies all secondary address prefixes.
summary	(Optional) Displays a brief list of the IS-IS topology.
level { **1** | **2** }	(Optional) Displays the IS-IS link-state topology for Level 1 or Level 2 independently.
multicast-intact	(Optional) Displays multicast intact information on the IS-IS topology.
systemid *system-id*	(Optional) Displays the information for the specified router only.
detail	(Optional) Displays detailed information on the IS-IS topology.

Command Default

No instance ID specified displays a list of connected routers in all areas for all the IS-IS instances.

Both Level 1 and Level 2 is configured if no level is specified.

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

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 37: 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 38: 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**] [**all***protocol*]

Syntax Description

afi-all	(Optional) Specifies all address families.
ipv4	(Optional) Specifies an IPv4 address family.
ipv6	(Optional) Specifies an IPv6 address family.
all	(Optional) Specifies all protocols for a given address family.
protocol	(Optional) Specifies a routing protocol. For the IPv4 address family, the options are: • **bgp** • **isis** • **ospf** • **rip** For the IPv6 address family, the options are: • **bgp** • **isis** • **ospfv3**

Command Default

If no address family is specified, the default is IPv4.

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

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 :

```
RP/0/RP0/CPU0:router# show protocols ipv4

IS-IS Router: uut
  System Id: 0000.0000.12a8
  IS Levels: level-1-2
  Manual area address(es):
    49.1515.1515
  Routing for area address(es):
    49.1515.1515
  Non-stop forwarding: Disabled
  Most recent startup mode: Cold Restart
  Topologies supported by IS-IS:
    IPv4 Unicast
      Level-1
        Metric style (generate/accept): Narrow/Narrow
        ISPF status: Disabled
      Level-2
        Metric style (generate/accept): Narrow/Narrow
        ISPF status: Disabled
      Redistributing:
        static
      Distance: 115
    IPv6 Unicast
      Level-1
        ISPF status: Disabled
      Level-2
        ISPF status: Disabled
      No protocols redistributed
      Distance: 45
  Interfaces supported by IS-IS:
    GigabitEthernet 0/6/0/0 is running actively (active in configuration)
```

This table describes the significant fields shown in the display.

Table 39: 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.

Field	Description
Non-stop forwarding	Status and name of NSF.
Most recent startup mode	Mode in which the most recent startup was performed.
Topologies supported by IS-IS	Address and subaddress family IS-IS are configured.
Metric style	Type, length, and value (TLV) objects accepted by IS-IS. To configure this value, see the metric-style wide, on page 115 command.
ISPF status	State of iSPF configuration for this IS-IS instance. Four states exist: • Disabled if iSPF has not been configured but is awaiting a full SPF to compile the topology for use by the iSPF algorithm. • Dormant if iSPF has been configured but is awaiting initial convergence before initializing. • Awake if iSPF has been configured but is awaiting a full SPF to compile the topology for use by the iSPF algorithm. • Active if IS-IS is ready to consider using the iSPF algorithm whenever a new route calculation needs to be run.
Redistributing	IS-IS is configured to redistribute IP static routes into Level 1 or Level 2. The **redistribute** command is used to configure redistribution.
Distance	Administrative distance.
Interfaces supported by IS-IS	Interfaces and their states currently supported by IS-IS. Both operational and configuration status are displayed.

The following example shows how to disable the IPv4 address family, with no output shown for IS-IS IPv4 instances from the **show protocols ipv4** command:

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router isis uut
RP/0/RP0/CPU0:router(config-isis)# no address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-isis)# commit

RP/0/RP0/CPU0:router# show protocols ipv4
```

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

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 7.0.12	This command was introduced.

Usage Guidelines

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

To enable the Simple Network Management Protocol (SNMP) server notifications (traps) available for IS-IS, use the **snmp-server traps isis** command in XR Config mode. To disable all available SNMP notifications, use the **no** form of this command.

snmp-server traps isis { **all** | *traps set* }

Syntax Description

all	Specifies all IS-IS SNMP server traps.
traps set	Specify any set of trap names.

Command Default

SNMP server traps notification is disabled.

Command Modes

XR Config mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

Task ID

Task ID	Operations
isis	read, write

Examples

The following examples show how to enable all SNMP server traps available for isis:

```
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
```

The following example shows how to enable area-mismatch lsp-error-detected trap:

```
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** }]

Syntax Description

initial-wait *initial*	Initial SPF calculation delay (in milliseconds) after a topology change. Range is 0 to 120000.
secondary-wait *secondary*	Hold time between the first and second SPF calculations (in milliseconds). Range is 0 to 120000.
maximum-wait *maximum*	Maximum interval (in milliseconds) between two consecutive SPF calculations. Range is 0 to 120000.
level { **1** | **2** }	(Optional) Enables the SPF interval configuration for Level 1 or Level 2 independently.

Command Default

initial-wait *initial* : 50 milliseconds

secondary-wait *secondary* : 200 milliseconds

maximum-wait *maximum* : 5000 milliseconds

Command Modes

Address family configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

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-interval ietf

To set an shortest path first (SPF) interval in IS-IS for SPF calculations, use the **spf-interval ietf** command in the System Admin Config mode. Use the **no** form of this command to enable the fabric bundle port.

spf-interval ietf [**initial-wait** *msec* | **short-wait** *msec* | **long-wait** *msec* | **learn-interval** *msec* | **holddown-interval** *msec*] [**level** { **1** | **2** }]

Syntax Description

spf-interval	Specifies the number of seconds between two consecutive SPF calculations.
ietf	Specifies Internet Engineering Task Force (IETF) RFC standard 8405.
initial-wait *msec*	Initial SPF calculation delay before running a route calculation. The initial-wait must be less than or equal to short-wait. Range is 0 to 120000. The default value is 50 milliseconds.
short-wait *msec*	Short SPF calculation delay before running a route calculation. The short-wait must be less than or equal to long-wait. Range is 0 to 120000. The default value is 200 milliseconds.
long-wait *msec*	Long SPF calculation delay before running a route calculation. Range is 0 to 120000. The default value is 5000 milliseconds.
learn-interval *msec*	Time To Learn interval for running a route calculation. The learn-interval must be less than or equal to holddown-interval. Range is 0 to 120000. The default value is 500 milliseconds.
holddown-interval *msec*	Hold-down interval for running a route calculation. Range is 0 to 120000. The default value is 10000 milliseconds.
level { **1** | **2** }	(Optional) Enables the SPF interval configuration for Level 1 or Level 2 independently.

Command Default

None

Command Modes

System Admin Config mode

Command History

Release	Modification
Release 7.7.1	This command was introduced.

Usage Guidelines

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

SPF calculations are performed only when the topology changes. They are not performed when external routes change.

Task ID

Task ID	Operations
is-is	read, write

Example

The following example shows how to configure IETF to set an SPF interval in IS-IS for SPF calculations.

```
Router# configure
Router(config)# router isis isp
Router(config-isis)# address-family ipv4 unicast
Router(config-isis-af)# spf-interval ietf?
initial-wait   Initial delay before running a route calculation [50]
short-wait     Short delay before running a route calculation [200]
long-wait      Long delay before running a route calculation [5000]
learn-interval Time To Learn interval for running a route calculation [500]
holddown-interval  Holddown interval for running a route calculation [10000]
level             Set SPF interval for one level only
Router(config-isis-af)# spf-interval ietf
Router(config-isis-af)# commit
```

The following **show** command displays the output with the new spf-interval algorithm. The output displays the actual delay taken to compute the SPF.

```
Router# show isis ipv4 spf-log last 5 detail
   IS-IS 1 Level 2 IPv4 Unicast Route Calculation Log
                    Time Total Trig.
Timestamp    Type    (ms) Nodes Count First Trigger LSP    Triggers
------------ ----- ----- ----- ----- -------------------- ------------------------
--- Wed Mar 16 2022 ---
15:31:49.763  FSPF     1     6     3         tb5-r4.00-00 LINKBAD PREFIXBAD
  Delay:                 101ms (since first trigger)
                         261177ms (since end of last calculation)
  Trigger Link:          tb5-r2.00
  Trigger Prefix:        34.1.24.0/24
  New LSP Arrivals:      0
  SR uloop:              No
  Next Wait Interval:    200ms
  RIB Batches:           1 (0 critical, 0 high, 0 medium, 1 low)
  Timings (ms):          +--Total--+
                          Real   CPU
    SPT Calculation:         1     1
    Route Update:            0     0
                         ----- -----
```

spf prefix-priority (IS-IS)

To assign a priority to an ISIS prefix for customizing the RIB update sequence, use the**spf prefix-priority** command in address family configuration mode. To restore default values, use the **no** form of this command.

spf prefix-priority [**level** { **1** | **2** }] { **critical** | **high** | **medium** } { *access-list-name* | **tag** *tag* }

Syntax Description

level { **1** | **2** }	(Optional) Enables the assignment of a priority to Level 1 or Level 2 independently.
critical	Assigns a critical priority.
high	Assigns a high priority.
medium	Assigns a medium priority.
access-list-name	Name of an access list.
tag *tag*	Specifies a tag to indicate priority. The *tag* argument range is 1 to 4294967295.

Command Default

By default, IPv4 prefixes with a length of 32 and IPv6 prefixes with a length of 128 are given medium priority. The remaining prefixes are given low priority.

Command Modes

Address family configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

Use the **spf prefix-priority** command to change the sequence of prefix updates to the RIB after an SPF is run. ISIS installs prefixes in the RIB according to the following priority order:

Critical > High > Medium > Low

The **spf prefix-priority** command supports prefix lists for the first three priorities. The unmatched prefixes are updated with low priority.

If a **spf prefix-priority** is specified, the default behavior of prioritizing either length 32 or 128 prefixes for IPv4 or IPv6, respectively, as **medium** is disabled.

Task ID

Task ID	Operations
isis	read, write

Examples

The following example shows how to set the prefix priorities:

```
RP/0/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 *address* argument must be in four-part, dotted-decimal notation.
/ *prefix-length*	Length of the IPv4 or IPv6 prefix. A decimal value that indicates how many of the high-order contiguous bits of the address compose the prefix (the network portion of the address). A slash must precede the decimal value.
ipv6-prefix	Summary prefix designated for a range of IPv6 prefixes. The *ipv6-prefix* argument must be in the form documented in RFC 2373, in which the address is specified in hexadecimal using 16-bit values between colons.
level { **1** | **2** }	(Optional) Redistributes routes into Level 1 or Level 2 and summarizes them with the configured address and mask value.
tag *tag*	Sets a tag value. The value range is 1- 4294967295.

Command Default

All redistributed routes are advertised individually.

Both Level 1 and Level 2 are configured if no level is specified.

Command Modes

Address family configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

Multiple groups of addresses can be summarized for a given level. Routes learned from other routing protocols can also be summarized. The metric used to advertise the summary is the smallest metric of all the more-specific routes. Use the **summary-prefix** command to help reduce the size of the routing table.

This command also reduces the size of the link-state packets (LSPs) and thus the link-state database. It also helps ensure stability, because a summary advertisement depends on many more specific routes. If one more-specific route flaps, in most cases, this flap does not cause a flap of the summary advertisement.

The drawback of summary addresses is that other routes might have less information to calculate the most optimal routing table for all individual destinations.



Note When IS-IS advertises a summary prefix, it automatically inserts the summary prefix into the IP routing table but labels it as a “discard” route entry. Any packet that matches the entry is discarded to prevent routing loops. When IS-IS stops advertising the summary prefix, the routing table entry is removed.

Task ID

Task ID	Operations
isis	read, write

Examples

The following example shows how to redistribute Open Shortest Path First (OSPF) routes into IS-IS:

The following example shows how to redistribute Open Shortest Path First (OSPF) routes into IS-IS. In the OSPF routing table, IPv6 routes exist for 3ffe:f000:0001:0000::/64, 3ffe:f000:0002:0000::/64, 3ffe:f000:0003:0000::/64, and so on. This example shows only 3ffe:f000::/24 advertised into IPv6 IS-IS Level 2.

```
RP/0/RP0/CPU0:router(config)# router isis isp
RP/0/RP0/CPU0:router(config-isis)# address-family ipv4 ipv6 unicast
RP/0/RP0/CPU0:router(config-isis-af)# redistribute ospf ospfv3 2 level-2
RP/0/RP0/CPU0:router(config-isis-af)# summary-prefix 10.10.10.10 3ffe:f000::/24 level-2
RP/0/RP0/CPU0:router(config-isis-af)# summary-prefix 10.10.10.10 3ffe:f000::/24 tag
```



OSPF Commands

- address-family (OSPF), on page 228
- adjacency stagger, on page 229
- area (OSPF), on page 231
- authentication (OSPF), on page 232
- authentication-key (OSPF), on page 234
- auto-cost (OSPF), on page 236
- capability opaque disable, on page 237
- clear ospf process, on page 238
- clear ospf redistribution, on page 239
- clear ospf routes, on page 240
- clear ospf statistics, on page 241
- cost (OSPF), on page 243
- cost-fallback (OSPF), on page 245
- database-filter all out (OSPF), on page 246
- dead-interval (OSPF), on page 247
- default-cost (OSPF), on page 249
- default-information originate (OSPF), on page 250
- default-metric (OSPF), on page 252
- demand-circuit (OSPF), on page 253
- disable-dn-bit-check, on page 255
- distance (OSPF), on page 256
- distance ospf, on page 258
- distribute-list, on page 259
- domain-id (OSPF), on page 261
- domain-tag, on page 263
- exchange-timer, on page 264
- fast-reroute (OSPFv2), on page 266
- fast-reroute per-link exclude interface, on page 268
- fast-reroute per-prefix exclude interface (OSPFv2), on page 269
- fast-reroute per-prefix lfa-candidate (OSPFv2), on page 270
- fast-reroute per-prefix remote-lfa (OSPFv2), on page 271
- fast-reroute per-prefix ti-lfa, on page 272
- fast-reroute per-prefix use-candidate-only (OSPFv2), on page 273

- flood-reduction (OSPF), on page 274
- hello-interval (OSPF), on page 275
- ignore lsa mospf, on page 277
- interface (OSPF), on page 278
- log adjacency changes (OSPF), on page 280
- loopback stub-network , on page 281
- max external-lsa, on page 282
- max-lsa, on page 284
- max-metric, on page 286
- maximum interfaces (OSPF), on page 289
- maximum paths (OSPF), on page 290
- maximum redistributed-prefixes (OSPF), on page 291
- message-digest-key, on page 293
- mpls ldp auto-config (OSPF), on page 296
- mpls ldp sync (OSPF), on page 297
- mpls traffic-eng (OSPF), on page 298
- mpls traffic-eng igp-intact (OSPF), on page 300
- mpls traffic-eng multicast-intact (OSPF), on page 301
- mpls traffic-eng ldp-sync-update (OSPF), on page 302
- mpls traffic-eng router-id (OSPF), on page 303
- mtu-ignore (OSPF), on page 305
- multi-area-interface, on page 307
- neighbor (OSPF), on page 309
- neighbor database-filter all out, on page 311
- network (OSPF), on page 312
- nsf (OSPF), on page 314
- nsf flush-delay-time (OSPF), on page 316
- nsf interval (OSPF), on page 317
- nsf lifetime (OSPF), on page 318
- nsr (OSPF), on page 319
- nssa (OSPF), on page 320
- ospf name-lookup, on page 321
- packet-size (OSPF), on page 322
- passive (OSPF), on page 324
- priority (OSPF), on page 326
- protocol shutdown, on page 328
- queue dispatch flush-lsa, on page 329
- queue dispatch incoming, on page 330
- queue dispatch rate-limited-lsa, on page 331
- queue dispatch spf-lsa-limit, on page 332
- queue limit, on page 333
- range (OSPF), on page 334
- redistribute (OSPF), on page 336
- retransmit-interval (OSPF), on page 340
- route-policy (OSPF), on page 342
- router-id (OSPF), on page 343

- router ospf (OSPF), on page 345
- security ttl (OSPF), on page 346
- segment-routing sr-prefer prefix-list, on page 348
- sham-link, on page 350
- show ospf, on page 351
- show ospf border-routers, on page 355
- show ospf database, on page 357
- show ospf flood-list, on page 372
- show ospf interface, on page 374
- show ospf mpls traffic-eng, on page 376
- show ospf message-queue, on page 381
- show ospf neighbor, on page 383
- show ospf request-list, on page 390
- show ospf retransmission-list, on page 392
- show ospf routes, on page 394
- show ospf sham-links, on page 399
- show ospf summary-prefix, on page 401
- show ospf virtual-links, on page 403
- show protocols (OSPF), on page 405
- snmp context (OSPF), on page 407
- snmp trap (OSPF), on page 409
- snmp trap rate-limit (OSPF), on page 410
- spf prefix-priority (OSPFv2), on page 411
- stub (OSPF), on page 413
- summary-prefix (OSPF), on page 414
- timers lsa group-pacing, on page 416
- timers lsa min-arrival, on page 417
- timers throttle lsa all (OSPF), on page 418
- timers throttle spf (OSPF), on page 420
- transmit-delay (OSPF), on page 422
- virtual-link (OSPF), on page 424
- vrf (OSPF), on page 426

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**]

Syntax Description

ipv4	Specifies IP Version 4 (IPv4) address prefixes.
unicast	(Optional) Specifies unicast address prefixes.

Command Default

An address family is not specified.

Command Modes

Router configuration

VRF configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

OSPF version 2 automatically provides routing services for IPv4 unicast topologies, so this command is redundant.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to configure the OSPF router process with IPv4 unicast address prefixes:

```
RP/0/RP0/CPU0:router(config)# router ospf 1
RP/0/RP0/CPU0:router(config-ospf)# address-family ipv4 unicast
```

adjacency stagger

To configure staggering of OSPF adjacency during reload, process restart, and process clear, use the **adjacency stagger** command in XR Config 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* }

Syntax Description

disable	Disables adjacency staggering.
initial-num-nbr	The initial number of simultaneous neighbors allowed to form adjacency to FULL in any area to bring up to FULL after a router reload, OSPF process restart, or OSPF process clear. Range is 1-65535. Default is 2.
max-num-nbr	The subsequent number of simultaneous neighbors allowed to form adjacency, per OSPF instance, after the initial set of OSPF neighbors have become FULL. Range is 1-65535. Default is 64.

Command Default

OSPF adjacency staggering is enabled.

Command Modes

XR Config mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

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*

Syntax Description

area-id	Identifier of an OSPF area. The *area-id* argument can be specified as either a decimal value or an IP address (dotted decimal) format. Range is 0 to 4294967295.

Command Default

No OSPF area is defined.

Command Modes

Router configuration

VRF configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

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 HundredGigE interface 0/2/0/0. HundredGigE 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 HundredGigE 0/2/0/0
```

authentication (OSPF)

To enable plain text, Message Digest 5 (MD5) authentication, or null authentication for an Open Shortest Path First (OSPF) interface, use the **authentication** command in the appropriate mode. To remove such authentication, use the **no** form of this command.

authentication [**message-digest** [**keychain** *keychain*] | **null**]
no authentication

Syntax Description

message-digest	(Optional) Specifies that MD5 is used.
keychain *keychain*	(Optional) Specifies a keychain name.
null	(Optional) Specifies that no authentication is used. Useful for overriding password or MD5 authentication if configured for an area.

Command Default

If this command is not specified in interface configuration mode, then the interface adopts the authentication parameter specified by the area.

If this command is not specified in area configuration mode, then the interface adopts the authentication parameter specified for the process.

If this command is not specified at any level, then the interface does not use authentication.

If no keyword is specified, plain text authentication is used.

Command Modes

Interface configuration

Area configuration

Router configuration

Virtual-link configuration

Multi-area interface configuration

Sham-link configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

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.

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
```

authentication-key (OSPF)

To assign a password to be used by neighboring routers that are using the Open Shortest Path First (OSPF) simple password authentication, use the **authentication-key** command in the appropriate mode. To remove a previously assigned OSPF password, use the **no** form of this command.

authentication-key [**clear** | **encrypted**] *password*
no authentication-key

Syntax Description

clear	(Optional) Specifies that the key be clear text.
encrypted	(Optional) Specifies that the key be encrypted using a two-way algorithm.
password	Any contiguous string up to 8 characters in length that can be entered from the keyboard. For example, *mypswd2*.

Command Default

If this command is not specified in interface configuration mode, then the interface adopts the OSPF password parameter specified by the area.

If this command is not specified in area configuration mode, then the interface adopts the OSPF password parameter specified for the process.

If this command is not specified at any level, then no password is specified.

Clear is the default if the **clear** or **encrypted** keyword is not specified.

Command Modes

Interface configuration

Area configuration

Router configuration

Virtual-link configuration

Multi-area configuration

Sham-link configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

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
```

auto-cost (OSPF)

To control how the Open Shortest Path First (OSPF) protocol calculates default metrics for the interface, use the **auto-cost** command in the appropriate mode. To revert to the default reference bandwidth, use the **no** form of this command.

auto-cost {**reference-bandwidth** *mbps* | **disable**}
no auto-cost {**reference-bandwidth** | **disable**}

Syntax Description

reference-bandwidth	*mbps*	Specifies a rate in Mbps (bandwidth). Range is 1 to 4294967.
disable		Assigns a cost based on interface type.

Command Default

mbps : 100 Mbps

Command Modes

Router configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

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
```

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

Command Default

Opaque LSAs are allowed.

Command Modes

XR Config mode

VRF configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

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.

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-name* [**vrf** { *vrf-name* | **all** }]] **process**

Syntax Description

process-name	(Optional) Name that uniquely identifies an OSPF routing process. The process name is defined by the **router ospf** command. If this argument is included, only the specified routing process is affected. Otherwise, all OSPF processes are reset.
vrf	(Optional) An OSPF VPN routing and forwarding (VRF) instance.
vrf-name	(Optional) Name of the OSPF VRF instance to be reset.
all	(Optional) Resets all OSPF VRF instances.

Command Default

No default behavior or value

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

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 343 command.

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
```

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 [*process-name* [**vrf** {*vrf-name* | **all**}]] **redistribution**

Syntax Description

process-name	(Optional) Name that uniquely identifies an OSPF routing process. The process name is defined by the **router ospf** command. If this argument is included, only the specified routing process is affected. Otherwise, all OSPF routes are cleared.
vrf	(Optional) OSPF VPN routing and forwarding (VRF) instance.
vrf-name	(Optional) Name of the OSPF VRF instance to be reset.
all	(Optional) Resets all OSPF VRF instances.

Command Default

No default behavior or value

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

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 [*process-name* [**vrf** {*vrf-name* | **all**}]] **routes**

Syntax Description

process-name	(Optional) Name that uniquely identifies an OSPF routing process. The process name is defined by the **router ospf** command. If this argument is included, only the specified routing process is affected. Otherwise, all OSPF routes are cleared.
vrf	(Optional) OSPF VPN routing and forwarding (VRF) instance.
vrf-name	(Optional) Name of the OSPF VRF instance to be reset.
all	(Optional) Resets all OSPF VRF instances.

Command Default

No default behavior or value

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to clear all OSPF routes from the OSPF routing table and recompute valid routes. When the OSPF routing table is cleared, OSPF routes in the global routing table are also recalculated.

```
RP/0/RP0/CPU0:router# clear ospf routes
```

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 [*process-name* [**vrf** {*vrf-name* | **all**}]] **statistics** [**neighbor** [*type interface-path-id*] [*ip-address*]]

Syntax Description

process-name	(Optional) Name that uniquely identifies an OSPF routing process. The process name is defined by the **router ospf** command. If this argument is included, only the specified routing process is affected. Otherwise, all OSPF statistics of neighbor state transitions are cleared.
vrf	(Optional) OSPF VPN routing and forwarding (VRF) instance.
vrf-name	(Optional) Name of the OSPF VRF instance to be reset.
all	(Optional) Resets all OSPF VRF instances.
neighbor	(Optional) Clears the state transition counters of the specified neighbor only.
type	(Optional) Interface type. For more information, use the question mark (?) online help function.
interface-path-id	(Optional) Physical interface or virtual interface. Use the **show interfaces** command to see a list of all interfaces currently configured on the router. For more information about the syntax for the router, use the question mark (?) online help function.
ip-address	(Optional) IP address of a specified neighbor for whom you want to clear the state transition counter.

Command Default

No default behavior or value

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

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
```

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*

Syntax Description

cost	Unsigned integer value expressed as the link-state metric. Range is 1 to 65535.

Command Default

If this command is not specified in interface configuration mode, then the interface adopts the cost parameter specified by the area.

If this command is not specified in area configuration mode, then the interface adopts the cost parameter specified for the process.

If this command is not specified at any level, then the cost is calculated by the **auto-cost** command.

Command Modes

Interface configuration

Area configuration

Router configuration

VRF configuration

Multi-area configuration

Sham-link configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

The link-state metric is advertised as the link cost in the router link advertisement. Cisco IOS XR software does not support type of service (ToS), so you can assign only one cost for each interface.

In general, the path cost is calculated using the following formula:

108 / bandwidth (the default auto cost is set to 100 Mbps)

This calculation is the default reference bandwidth used by the auto-costing calculation which establishes the interface auto-cost The **auto-cost** command can set this reference bandwidth to some other value. The **cost** command is used to override the auto-costing calculated default value for interfaces.

Using this formula, the default path cost is 1 for any interface that has a link bandwidth of 100 Mbps or higher. If this value does not suit the network, configure the reference bandwidth for auto calculating costs based on the link bandwidth.

The value set by the **cost** command overrides the cost resulting from the **auto-cost (OSPF)** command.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to set the cost value to 65 for HundredGigE 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 HundredGigE 0/1/0/1
RP/0/RP0/CPU0:router(config-ospf-ar-if)# cost 65
```

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

Syntax Description

cost **threshold**	Unsigned integer value expressed as the link-state metric. Range is 1 to 65535, but typically, cost-fallback value is supposed to be set to a value higher than the normal cost.
bandwidth	Unsigned integer value expressed in Mbits per second. Range is 1 to 4294967.

Command Default

If this command is not specified in interface configuration mode, the currently effective interface cost takes effect even when the cumulative bandwidth goes down below the maximum bandwidth. Unlike the interface cost command, this cost-fallback command is available only under interface configuration mode; it is not available in area or process level. Unlike other interface specific parameters, no inheritance will take place from area or process level if this command is not specified at interface level.

Command Modes

Interface configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

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:

The following example shows how to set the cost-fallback value for Bundle-Ether:

```
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-Ether
RP/0/RP0/CPU0:router(config-ospf-ar-if)# cost-fallback 1000 threshold 300
```

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 7.0.12	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 database-filter all out (OSPF), on page 246 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 HundredGigE 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 HundredGigE 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*

Syntax Description

seconds	Integer that specifies the interval (in seconds). Range is 1 to 65535. The value must be the same for all nodes on the network.

Command Default

If this command is not specified in interface configuration mode, then the interface adopts the dead interval parameter specified by the area.

If this command is not specified in area configuration mode, then the interface adopts the dead interval parameter specified for the process.

If this command is not specified at any level, then the dead interval is four times the interval set by the **hello-interval** (OSPF) command.

Command Modes

Interface configuration

Area configuration

Router configuration

Virtual-link configuration

VRF configuration

Multi-area configuration

Sham-link configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

The dead interval value must be the same for all routers and access servers on a specific network.

If the hello interval is configured, the dead interval value must be larger than the hello interval value. The dead interval value is usually configured four times larger than the hello interval value.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to set the OSPF dead interval to 40 seconds:

```
RP/0/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 HundredGigE 0/1/0/1
RP/0/RP0/CPU0:router(config-ospf-ar-if)# dead-interval 40
```

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*

Syntax Description

cost	Cost for the default summary route used for a stub or NSSA area. The acceptable value is a 24-bit number.

Command Default

cost : 1

Command Modes

Area configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

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 HundredGigE 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 HundredGigE 0/4/0/3
```

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*]

Syntax Description

always	(Optional) Always advertises the default route regardless of whether the routing table has a default route.
metric *metric-value*	(Optional) Specifies the metric used for generating the default route. The default metric value is 1. Range is 1 to 16777214.
metric-type *type-value*	(Optional) Specifies the external link type associated with the default route advertised into the OSPF routing domain. It can be one of the following values: **1**—Type 1 external route **2**—Type 2 external route
route-policy *policy-name*	(Optional) Specifies that a routing policy be used and the routing policy name.

Command Default

When you do not use this command in router configuration mode, no default external route is generated into an OSPF routing domain.

metric-value : 1

type-value : 2

Command Modes

Router configuration

VRF configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

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 8000 Series Routers*.

For information about routing policies, see the *Routing Policy Commands* chapter in the *Routing Command Reference for Cisco 8000 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 igrp 108 metric 100
RP/0/RP0/CPU0:router(config-ospf)#default-information originate metric 100 metric-type 1
```

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*

Syntax Description

value	Default metric value appropriate for the specified routing protocol. Range is 1 to 16777214.

Command Default

Built-in, automatic metric translations, as appropriate for each routing protocol.

Command Modes

Router configuration

VRF configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

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
```

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**]

Syntax Description

disable	(Optional) Disables the interface as an OSPF demand circuit.
enable	(Optional) Enables the interface as an OSPF demand circuit.

Command Default

If this command is not specified in interface configuration mode, then the interface adopts the demand circuit parameter specified by the area.

If this command is not specified in area configuration mode, then the interface adopts the demand circuit parameter specified for the process.

If this command is not specified at any level, then the circuit is not a demand circuit.

Command Modes

Interface configuration

Area configuration

Router configuration

VRF configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

On point-to-point interfaces, only one end of the demand circuit must be configured with this command. Periodic hello messages are suppressed and periodic refreshes of link-state advertisements (LSAs) do not flood the demand circuit. Use the **demand-circuit** command to allow the underlying data link layer to be closed when the topology is stable. In point-to-multipoint topology, only the multipoint end must be configured with this command.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to set the configuration for an OSPF demand circuit:

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router ospf 1
RP/0/RP0/CPU0:router(config-ospf)# demand-circuit
```

disable-dn-bit-check

To specify that down bits should be ignored, use the **disable-dn-bit-check** command in VPN routing and forwarding (VRF) configuration mode. To specify that down bits should be considered, use the **no** form of this command.

disable-dn-bit-check

Command Default

Down bits are considered.

Command Modes

VRF configuration mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to specify that down bits be ignored:

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router ospf 1
RP/0/RP0/CPU0:router(config-ospf)# vrf v1
RP/0/RP0/CPU0:router(config-ospf-vrf)# disable-dn-bit-check
```

distance (OSPF)

To define an administrative distance, use the **distance** command in an appropriate configuration mode. To remove the **distance** command from the configuration file and restore the system to its default condition in which the software removes a distance definition, use the **no** form of this command.

distance *weight* [*ip-address* *wildcard-mask* [*access-list-name*]]

Syntax Description

weight	Administrative distance. Range is 10 to 255. Used alone, the *weight* argument specifies a default administrative distance that the software uses when no other specification exists for a routing information source. Routes with a distance of 255 are not installed in the routing table. lists the default administrative distances. Table 1
ip-address	(Optional) IP address in four-part, dotted-decimal notation.
wildcard-mask	(Optional) Wildcard mask in four-part, dotted decimal format. A bit set to 1 in the *mask* argument instructs the software to ignore the corresponding bit in the address value.
access-list-name	(Optional) Name of an IP access list to be applied to incoming routing updates.

Command Default

If this command is not specified, then the administrative distance is the default, as specified in Table 1.

Command Modes

Router configuration

VRF configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

An administrative distance is an integer from 10 to 255. In general, the higher the value, the lower the trust rating. An administrative distance of 255 means that the routing information source cannot be trusted at all and should be ignored. Weight values are subjective; no quantitative method exists for choosing weight values.

If an access list is used with this command, it is applied when a network is being inserted into the routing table. This behavior allows you to filter networks based on the IP prefix supplying the routing information. For example, you could filter possibly incorrect routing information from networking devices not under your administrative control.

The order in which you enter **distance** commands can affect the assigned administrative distances in unexpected ways (see the “Examples” section for further clarification).

This table lists default administrative distances.

Table 40: Default Administrative Distances

Route Source	Default Distance
Connected interface	0

Route Source	Default Distance
Static route out on interface	0
State route to next-hop	1
External BGP	20
OSPF	110
IS-IS	115
RIP version 1 and 2	120
Internal BGP	200
Unknown	255

Task ID

Task ID	Operations
ospf	read, write

Examples

In the following example, the **router ospf** command sets up OSPF routing instance1. The first **distance** command sets the default administrative distance to 255, which instructs the software to ignore all routing updates from networking devices for which an explicit distance has not been set. The second **distance** command sets the administrative distance for all networking devices on the Class C network 192.168.40.0 0.0.0.255 to 90.

```
RP/0/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
```

distance ospf

To define Open Shortest Path First (OSPF) route administrative distances based on route type, use the **distance ospf** command in router configuration mode. To restore the default value, use the **no** form of this command.

distance ospf { **intra-area** | **inter-area** | **external** } *distance*
no distance ospf

Syntax Description

intra-area | **inter-area** | **external**	Sets the type of area. It can be one of the following values: **intra-area** —All routes within an area. **inter-area** —All routes from one area to another area. **external** —All routes from other routing domains, learned by redistribution. Any combination of the above areas is allowed.
distance	Route administrative distance.

Command Default

distance : 110

Command Modes

Router configuration

Usage Guidelines

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:

```
Router# configure
Router(config)# router ospf 1
Router(config-ospf)# redistribute ospf 2
Router(config-ospf)# distance ospf external 200
Router(config-ospf)# exit
Router(config)# router ospf 2
Router(config-ospf)# redistribute ospf 1
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**] } | **route-policy** *route-policy-name* **in** }

Syntax Description

access-list-name	Standard IP access list name. The list defines which networks are to be received and which are to be suppressed in routing updates.
in	Applies the access list or route-policy to incoming routing updates.
out	Applies the access list to outgoing routing updates. The **out** keyword is available only in router configuration mode.
bgp	(Optional) Applies the access list to BGP routes.
connected	(Optional) Applies the access list to connected routes.
ospf	(Optional) Applies the access list to OSPF routes (not the current OSPF process).
static	(Optional) Applies the access list to statically configured routes.
route-policy *route-policy-name*	Specifies the route-policy to filter OSPF prefixes.

Command Default

If this command is not specified in interface configuration mode, then the interface adopts the distribute list parameter specified by the area.

If this command is not specified in area configuration mode, then the interface adopts the distribute list parameter specified for the process.

If this command is not specified at any level, then the distribute list is disabled.

Command Modes

Interface configuration

Area configuration

Router configuration

VRF configuration

Multi-area configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

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.

Use the **route-policy** *route-policy-name* keyword and argument to allow use of route policies to filter OSPF prefixes.



Note Either an access-list, or a route-policy can be used in a single command, not both. Configuring the command with access-list removes the route-policy configuration, and vice versa.

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 HundredGigE 0/1/0/3
```

domain-id (OSPF)

To specify the Open Shortest Path First (OSPF) VPN routing and forwarding (VRF) domain ID, use the **domain-id** command in VRF configuration mode. To remove an OSPF VRF domain ID, use the **no** form of this command.

domain-id [**secondary**] **type** [**0005** | **0105** | **0205** | **8005**] **value** *value*

Syntax Description

secondary		(Optional) OSPF secondary domain ID.
type		Primary OSPF domain ID in hex format.
value	*value*	OSPF domain ID value in hex format (six octets).

Command Default

No domain ID is specified.

Command Modes

VRF configuration mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

An OSPF domain id must be explicitly configured. The OSPF domain ID helps OSPF determine how to translate a prefix received through Border Gateway Protocol (BGP) from the remote provider edge (PE). If the domain IDs match, OSPF generates a Type 3 link state advertisement (LSA). If the domain IDs do not match, OSPF generates a Type 5 LSA.

There is only one primary domain ID. There can be multiple secondary domain IDs.

Note When an IOS XR router and an IOS router are configured as peers, the two Domain IDs must match. Manually configure the IOS XR Domain ID value to match the IOS default Domain ID value. This ensures that the routes have route code "OIA" because they are learned as inter-area routes. If the Domain IDs do not match, the routes have route code, "O-E2" because they are learned as external routes. Use the **show ip ospf** command to get the OSPF Domain ID from the IOS router. Then, set the IOS XR Domain ID to the same value using the **domain-id** command.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to specify a domain ID:

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

```
RP/0/RP0/CPU0:router(config)# router ospf o1
RP/0/RP0/CPU0:router(config-ospf)# vrf v1
RP/0/RP0/CPU0:router(config-ospf-vrf)# domain-id type 0105 value AABBCCDDEEFF
```

domain-tag

To specify the Open Shortest Path First (OSPF) VPN routing and forwarding (VRF) domain tag, use the **domain-tag** command in VRF configuration mode. To remove an OSPF VRF domain tag, use the **no** form of this command.

domain-tag *tag*

Syntax Description

tag	OSPF domain tag as a 32-bit value. The valid range is 0 to 4294967295.

Command Default

No OSPF VRF domain tag is specified.

Command Modes

VRF configuration mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

The domain tag is added to any Type 5 link state advertisements (LSAs) generated as a result of VPN-IP routes received from Border Gateway Protocol (BGP). The domain-tag is derived from BGP autonomous system number (ASN).

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to specify the domain tag:

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router ospf o1
RP/0/RP0/CPU0:router(config-ospf)# domain-tag 234
```

exchange-timer

To set a time limit to prevent an indefinite EXCHANGE or LOADING phase during adjacency formation, use the **exchange-timer** command in router configuration mode.

exchange-timer *exchange-time-in-minutes* **hold-time** *hold-time-in-minutes* **recovery-count** *recovery-count*

Syntax Description

exchange-time-in-minutes	Duration in minutes, after which the interface is brought down if adjacencies remain in the exchange phase, in minutes. The range is from 1 to 35791394. The default value is 10.
hold-time	Specifies the option to set the time after which a recovery attempt will be made.
hold-time-in-minutes	Duration in minutes until which interface will be kept in the down state before attempting recovery, in minutes. The range is from 1 to 35791394. The default value is 5 minutes.
recovery-count	Specifies the number of attempts the router will make to bring up the adjacency towards the ASBR before keeping the adjacency shut indefinitely.
recovery-count	The number of attempts the router will make to recover the interface before the interface is permanently disabled. The recovery count value resets to the default value of 300 when the interface adjacency reaches the FULL state with the neighbor. The range is from 1 to 4294967294. The default value is 300.

Command Default

The exchange timer operates using its default values.

Command Modes

Router Configuration

Command History

Release	Modification
Release 7.10.1	This command was introduced.

Usage Guidelines

To adjust the default parameters of the exchange timer, you can use the **exchange-timer** command. We recommend configuring this command with values that suit your specific scale requirements

By default, the exchange timer is enabled with the following default values:

- Default Exchange Time: 10 minutes
- Default Hold Time: 5 minutes
- Default Recovery Count: 300

These default values have been chosen to ensure that the interface remains operational for an extended period before being permanently brought down. The total duration for attempting recovery is calculated as follows:

```
(Exchange Time + Hold Time) * Recovery Count = 3.125 days
```

Task ID

Task ID	Operation
OSPF	read, write

Example

The following example shows how to set a time limit to prevent an indefinite EXCHANGE or LOADING phase during adjacency formation.

```
Router(config)# router ospf 1
Router(config-ospf)# exchange-timer 357913 hold-time 35791 recovery-count 4294967294
```

fast-reroute (OSPFv2)

To enable IP fast reroute loop-free alternate (LFA) computation, use the **fast-reroute** command in the appropriate OSPF configuration mode. To disable the IP fast reroute loop-free alternate computation, use the **no** form of this command.

To disable loop-free alternate computation that is enabled on a higher level, use the **fast-reroute** command with **disable** keyword.

fast-reroute { **per-link** | **per-prefix** } [**disable**]

Syntax Description

per-link	Enables per-link loop-free alternate computation.
per-prefix	Enables per-prefix loop-free alternate computation.
disable	(Optional) Disables loop-free alternate computation that was enabled on a higher level.

Command Default

IP fast-reroute LFA computation is disabled.

Command Modes

Area configuration

Interface configuration

Router configuration

VRF configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

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 HundredGigE 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 HundredGigE 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*

Syntax Description

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

Command Default

No interfaces are excluded.

Command Modes

Interface configuration

Area configuration

Router configuration

VRF configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

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 HundredGigE
 0/3/2/1
```

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*

Syntax Description

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

Command Default

No interfaces are excluded.

Command Modes

Interface configuration

Area configuration

Router configuration

VRF configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

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 HundredGigE
0/6/0/1
```

fast-reroute per-prefix lfa-candidate (OSPFv2)

To add interfaces to the LFA candidate list, use the **fast-reroute per-prefix lfa-candidate** command in interface configuration mode. To disable this feature, use the **no** form of this command.

fast-reroute per-prefix lfa-candidate [*interface-name*]
no fast-reroute per-prefix lfa-candidate [*interface-name*]

Syntax Description

interface-name	Specifies name of the interface to add to the LFA candidate list.

Command Default

No interfaces are added to the candidate list.

Command Modes

Interface configuration

Area configuration

Router configuration

Usage Guidelines

General ios xr variables to use this command, you must be in a user group associated with a task group that includes the proper task IDs. If you suspect user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

This example shows how to add an interface to LFA candidates:

```
Router#configure
Router(config)#router ospf 100
Router(config-ospf)#fast-reroute per-prefix lfa-candidate interface GigabitEthernet 0/6/0/0
```

fast-reroute per-prefix remote-lfa (OSPFv2)

To configure fast-reroute per-prefix remote loop-free alternate (rLFA) computation for an OSPFv2 process, use the **fast-reroute per-prefix remote-lfa** command in the appropriate OSPF configuration mode. To disable this feature, use the **no** form of this command.

fast-reroute per-prefix remote-lfa { **disable** | **maximum-cost** *path-cost* | **tunnel mpls-ldp** }
no fast-reroute per-prefix remote-lfa

Syntax Description

remote-lfa	Enables remote LFA backup computation
maximum-cost *path-cost*	Sets the cost option to limit the range of remote LFAs. Range for path-cost is 1 to 4294967295.
tunnel mpls-ldp	Enables remote LFA computation using tunnel interfaces.
disable	Selectively disables remote LFA calculation under one or more areas.

Command Default

Remote LFA FRR computation is disabled.

Command Modes

Router configuration

Area configuration

Interface configuration

Usage Guidelines

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

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.

This example shows how to enable fast-reroute per-prefix remote LFA computation for tunnel interfaces:

```
Router(config)#router ospf 1
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:

```
Router(config)#router ospf 1
Router(config-ospf)#fast-reroute per-prefix remote-lfa maximum-cost 2
```

Related Commands

Command	Description
Fast reroute ospfv2	

fast-reroute per-prefix ti-lfa

To enable per-prefix Topology Independent Loop-Free Alternate Fast Reroute (TI-LFAFRR) computation to reroute traffic around link failures, use the **fast-reroute per-prefix ti-lfa** command in the required configuration mode.

fast-reroute per-prefix ti-lfa { **enable** | **disable** }

Syntax Description

enable	Enables per-prefix Topology Independent Loop-Free Alternate Fast Reroute (TI-LFAFRR) computation to reroute traffic around link failures.
disable	Disables per-prefix Topology Independent Loop-Free Alternate Fast Reroute (TI-LFAFRR) computation.

Command Default

TI-LFAFRR computation is not enabled by default.

Command Modes

Interface configuration

Area configuration

Router configuration

Command History

Release	Modification
Release 5.3.0	This command was introduced.

Usage Guidelines

OSPFv2 Loop-Free Alternate Fast Reroute TI-LFA is built on top of per-prefix LFA and as such requires per-prefix LFA to be enabled on the interface, where TI-LFA is enabled.

Task ID

Task ID	Operation
ospf	read, write

This example shows how to enable per-prefix Topology Independent Loop-Free Alternate Fast Reroute computation for the interface:

```
Router#configure
Router(config)#router ospf
Router(config-ospf)#area 1
Router(config-ospf-ar)#interface
GigabitEternet0/5/0/0 fast-reroute per-prefix ti-lfa enable
```

fast-reroute per-prefix use-candidate-only (OSPFv2)

To restrict the backup interfaces to those that are present on the LFA candidate list, use the **fast-reroute per-prefix use-candidate-only** command in router OSPF configuration mode. To disable this feature, use the **no** form of this command.

fast-reroute per-prefix use-candidate-only [**enable** | **disable**]
fast-reroute per-prefix use-candidate-only

Syntax Description

enable	Enables backup selection from candidate-list only.
disable	Disables backup selection from candidate-list only.

Command Default

Disabled.

Command Modes

Router OSPF configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

Task ID

Task ID	Operation
ospf	read, write

This example shows how to restrict the backup interfaces to those that are present on the LFA candidate list:

```
RP/0/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**]

Syntax Description

enable	(Optional) Turns on this functionality at a specific level.
disable	(Optional) Turns off this functionality at a specific level.

Command Default

If this command is not specified in interface configuration mode, then the interface adopts the flood reduction parameter specified by the area.

If this command is not specified in area configuration mode, then the interface adopts the flood reduction parameter specified for the process.

If this command is not specified at any level, then flood reduction is disabled.

Command Modes

Interface configuration

Area configuration

Router configuration

VRF configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

All routers supporting the OSPF demand circuit are compatible and can interact with routers supporting flooding reduction.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to reduce the flooding of unnecessary LSAs for area 0:

```
RP/0/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 HundredGigE 0/1/0/3
RP/0/RP0/CPU0:router(config-ospf-ar-if)# flood-reduction
```

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*

Syntax Description

seconds	Interval (in seconds). The value must be the same for all nodes on a specific network. Range is 1 to 65535.

Command Default

If this command is not specified in interface configuration mode, then the interface adopts the hello interval parameter specified by the area.

If this command is not specified in area configuration mode, then the interface adopts the hello interval parameter specified for the process.

If this command is not specified at any level, then the hello interval is 10 seconds (broadcast) or 30 seconds (non-broadcast).

Command Modes

Interface configuration

Area configuration

Router configuration

Virtual-link configuration

VRF configuration

Multi-area configuration

Sham-link configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

The hello interval value is advertised in the hello packets. The shorter the hello interval, the faster topological changes are detected, but more routing traffic occurs. This value must be the same for all routers and access servers on a specific network.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to set the interval between hello packets to 15 seconds:

```
RP/0/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 HundredGigE 0/1/0/1
RP/0/RP0/CPU0:router(config-ospf-ar-if)# hello-interval 15
```

ignore lsa mospf

To suppress the sending of syslog messages when the router receives link-state advertisement (LSA) Type 6 multicast Open Shortest Path First (MOSPF) packets, which are unsupported, use the **ignore lsa mospf** command in an appropriate configuration mode. To restore the sending of syslog messages, use the **no** form of this command.

ignore lsa mospf

This command has no keywords or arguments.

Command Default

When you do not specify this command in router configuration mode, each MOSPF packet received by the router causes the router to send a syslog message.

Command Modes

Router configuration

VRF configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

Cisco routers do not support LSA Type 6 (MOSPF), and they generate syslog messages if they receive such packets. If the router is receiving many MOSPF packets, you might want to configure the router to ignore the packets and thus prevent a large number of syslog messages.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to configure the router to suppress the sending of syslog messages when it receives MOSPF packets:

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router ospf 1
RP/0/RP0/CPU0:router(config-ospf)# ignore lsa mospf
```

interface (OSPF)

To define the interfaces on which the Open Shortest Path First (OSPF) protocol runs, use the **interface** command in area configuration mode. To disable OSPF routing for interfaces, use the **interface** form of this command.

interface *type* *interface-path-id*

Syntax Description

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

Command Default

When you do not specify this command in configuration mode, OSPF routing for interfaces is not enabled.

Command Modes

Area configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

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 HundredGigE 4/0/0/3
!
RP/0/RP0/CPU0:router(config-ospf)# area 2
RP/0/RP0/CPU0:router(config-ospf-ar)# interface HundredGigE 0/1/0/3
!
RP/0/RP0/CPU0:router(config-ospf)# area 3
RP/0/RP0/CPU0:router(config-ospf-ar)# interface HundredGigE 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 HundredGigE 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 XR Config 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

XR Config mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

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 stubnetworks, use the **no** form of this command.

loopback stub-network [**enable** | **disable**]

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 7.0.12	This command was introduced.

Usage Guidelines

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
```

max external-lsa

To specify the maximum number of external LSAs that the router will accept from a neighbor, use the **max-external-lsa** command in router configuration mode.

max-external-lsa *lsa-limit* *logging-threshold-percentage* [**suppress-neighbor** | **warning-only**]

Syntax Description

lsa-limit	Specifies the maximum number of external LSAs that a router will accept from a neighbor. If the neighbor sends more than the specified number of LSAs, the adjacency will be brought down and kept down for a specified interval. The range is from 1 to 4294967294. The default value is 50,000.
logging-threshold-percentage	Specifies the percentage of the total number of external LSAs from an ASBR at which a warning log is given. The default value is 75.
suppress-neighbor	Specifies that the link towards an ASBR be brought down if the ASBR is the immediate neighbor. This option is enabled by default.
warning-only	Specifies the option to overwrite the suppress-neighbor behavior. If the **warning-only** keyword is configured instead of **suppress-neighbor** keyword, the link to the ASBR is not brought down when the limit of external LSAs is reached. Only alert logs is given.

Command Default

If the command is not configured, the ASBR does not monitor unsuccessful adjacencies towards other ASBRs caused by the ASBR reaching the configured maximum limit for external Link State Advertisements (LSAs). Additionally, the ASBR does not incorporate a counter that, upon reaching the specified number of attempts to establish adjacency, causes the neighbor to cease further efforts, triggering permanent adjacency downtime.

Command Modes

Router Configuration

Command History

Release	Modification
Release 7.10.1	This command was introduced.

Usage Guidelines

We recommend configuring this command with values that suit your specific scale requirements.

Task ID

Task ID	Operation
OSPF	read, write

Example

The following example show how to configure the maximum number of external LSAs, the warning log percentage, and the link-down option for an OSPF router.

```
Router(config)# router ospf 1
Router(config-ospf)# area 0
Router(config-ospf-ar)# max-external-lsa 50000 suppress-neighbor
```

max-lsa

To limit the number of nonself-generated link-state advertisements (LSAs) that an Open Shortest Path First (OSPF) routing process can keep in the OSPF link-state database (LSDB), use the **max-lsa** command in XR Config 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*]

Syntax Description

max	Maximum number of nonself-generated LSAs the OSPF process can keep in the OSPF LSBD.
threshold	(Optional) The percentage of the maximum LSA number, as specified by the maximum-number argument, at which a warning message is logged. The default is 75 percent.
warning-only	(Optional) Specifies that only a warning message is sent when the maximum limit for LSAs is exceeded. Disabled by default.
ignore-time *value*	(Optional) Specifies the time, in minutes, to ignore all neighbors after the maximum limit of LSAs has been exceeded. The default is 5 minutes.
ignore-count *value*	(Optional) Specifies the number of times the OSPF process can consecutively be placed into the ignore state. The default is 5 times.
reset-time *value*	(Optional) Specifies the time, in minutes, after which the ignore count is reset to zero. The default is 2 times **ignore-time** .

Command Default

Disabled

Command Modes

XR Config mode

VRF configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

This command allows you to protect the OSPF routing process from the large number of received LSAs that can result from a misconfiguration on another router in the OSPF domain (for example, the redistribution of a large number of IP prefixes to OSPF).

When this feature is enabled, the router keeps count of the number of all received (nonself-generated) LSAs. When the configured *threshold* value is reached, an error message is logged. When the configured *max* number of received LSAs is exceeded, the router stops accepting new LSAs.

If the count of received LSAs is higher than the configured *max* number after one minute, the OSPF process disables all adjacencies in the given context and clears the OSPF database. This state is called the ignore state. In this state, all OSPF packets received on all interfaces belonging to the OSPF instance are ignored and no OSPF packets are generated on its interfaces. The OSPF process remains in the ignore state for the duration

of the configured **ignore-time** . When the **ignore-time** expires, the OSPF process returns to normal operation and starts building adjacencies on all its interfaces.

To prevent the OSPF instance from endlessly oscillating between its normal state and the ignore state, as a result of the LSA count immediately exceeding the *max* number again after it returns from the ignore state, the OSPF instance keeps a count of how many times it has been in the ignore state. This counter is called the **ignore-count** . If the **ignore-count** exceeds its configured value, the OSPF instance remains in the ignore state permanently.

To return the OSPF instance to its normal state, you must issue the **clear ip ospf** command. The **ignore-count** is reset to zero if the LSA count does not exceed the *max* number again during the time configured by the **reset-time** keyword.

If you use the **warning-only** keyword, the OSPF instance never enters the ignore state. When LSA count exceeds the *max* number, the OSPF process logs an error message and the OSPF instance continues in its normal state operation.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to configure the OSPF instance to accept 12000 nonself-generated LSAs in the global routing table, and 1000 nonself-generated LSAs in VRF V1.

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router ospf 0
RP/0/RP0/CPU0:router(config-ospf)# max-lsa 12000
RP/0/RP0/CPU0:router(config-ospf)# vrf V1
RP/0/RP0/CPU0:router(config-ospf)# max-lsa 1000
```

The following example shows how to display the current status of the OSPF instance:

```
RP/0/RP0/CPU0:router# show ospf 0

  Routing Process "ospf 0" with ID 10.0.0.2
  NSR (Non-stop routing) is Disabled
  Supports only single TOS(TOS0) routes
  Supports opaque LSA
  It is an area border router
  Maximum number of non self-generated LSA allowed 12000
     Current number of non self-generated LSA 1
     Threshold for warning message 75%
     Ignore-time 5 minutes, reset-time 10 minutes
     Ignore-count allowed 5, current ignore-count 0
```

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 XR Config 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**]

Syntax Description

router-lsa	Always originates router link-state advertisements (LSAs) with the maximum metric.
external-lsa *overriding metric*	(Optional) Overrides the external-lsa metric with the max-metric value.The *overriding metric* argument specifies the number of in-summary-LSAs. The range is 1 to 16777215>. The default is 16711680.
include-stub	(Optional) Advertises stub links in router-LSA with the max-metric value (0xFFFF).
on-proc-migration *time*	(Optional) Sets the maximum metric temporarily after a process migration to originate router-LSAs with the max-metric value. The *time* range is 5 to 86400 seconds.
on-proc-restart *time*	(Optional) Sets the maximum metric temporarily after a process restart to originate router-LSAs with the max-metric value. The *time* range is 5 to 86400 seconds.
on-startup *time*	(Optional) Sets the maximum metric temporarily after a reboot to originate router-LSAs with the max-metric value. The *time* range is 5 to 86400 seconds.
on-switchover *time*	(Optional) Sets the maximum metric temporarily after a switchover to originate router-LSAs with the max-metric value. The *time* range is 5 to 86400 seconds. **Note** OSPF will not populate maximum metric on the router's generated LSAs, when the OSPF routing process is configured to support Nonstop Routing (NSR) or Nonstop Forwarding/Graceful restart (NSF/GR).
wait-for-bgp	(Optional) Causes OSPF to originate router LSAs with the maximum metric and allows Border Gateway Protocol (BGP) to decide when to start originating router LSAs with a normal metric instead of the maximum metric.
summary-lsa	(Optional) specifies the number of in summary-LSAs. The range is 1 to 16777215. The default is 16711680.

Command Default

Router LSAs are originated with normal link metrics.

overriding-metric :16711680

Command Modes

XR Config mode

VRF configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

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*

Syntax Description

number-interfaces	Number of interfaces. Range is 1 to 4294967295.

Command Default

If the command is not specified, the default is 1024.

Command Modes

Router configuration

VRF configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

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
```

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*

Syntax Description

maximum-routes-number	Maximum number of parallel routes that OSPF can install in a routing table. Range is 1 to 64. **Note** The maximum number of paths that can be configured is 64.

Command Default

The default value for maximum-paths depends on the platform supported maximum-paths value. 32 paths

Command Modes

Router configuration

VRF configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

The default value for maximum-path depends on the platform supported maximum-path value. Any custom value you define for the maximum-path parameter must be within the maximum value supported by platform. The configuration will be rejected if the value you have specified is more then what the platform supports.

When the maximum number of parallel routes is reduced, all existing paths are pruned and paths reinstalled at the new maximum number. During this route-reduction period, you may experience some packet loss for a few seconds. This may impact route traffic.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to allow a maximum of two paths to a destination:

```
RP/0/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**]

Syntax Description

maximum	Number of routes. Range is 1 to 4294967295.
threshold-value	(Optional) Threshold value (as a percentage) at which to generate a warning message. Range is 1 to 100.
warning-only	(Optional) Gives only a warning when the limit is exceeded.

Command Default

If the command is not specified, the default is 10000.

The threshold value defaults to 75 percent.

Command Modes

Router configuration

VRF configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.
Release 7.9.1	• Support for the maximum redistributed-prefixes limit was added for the prefixes redistributed as Type-3 LSAs. • Support for counting only the number of redistributed prefixes instead of the number of LSAs was added. (If multiple prefixes are aggregated in an LSA using the summary-prefix command, individual prefixes will be counted)

Usage Guidelines

- Use the **maximum redistributed-prefixes** command to increase or decrease the maximum number of prefixes (also referred to as routes) redistributed for an OSPF process.

  If the maximum value specified is less than the existing number of redistributed prefixes, then the prefixes are reset and the redistribution is restarted.

- Starting Cisco IOS XR Software Release 7.9.1, when the **summary-prefix** command is used to create aggregate LSAs for the routes redistributed into OSPF, each individual redistributed prefix mapped under the subnet advertised in the LSA is counted instead of the number of LSAs. For example, if 5 redistributed prefixes are aggregated by one LSA using the **summary-prefix** command, only the number of redistributed prefixes, which is 5 in this case, is counted and not the number of LSAs, which is 1.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to configure a maximum number of routes that can be redistributed for an OSPF routing process:

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router ospf 109
RP/0/RP0/CPU0:router(config-ospf)# maximum redistributed-prefixes 15000
```

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* }

Syntax Description

key-id	Key number. Range is 1 to 255.
md5	Enables OSPF MD5 authentication.
key	Alphanumeric string of up to 16 characters.
clear	Specifies that the key be clear text.
encrypted	Specifies that the key be encrypted using a two-way algorithm.

Command Default

If this command is not specified in interface configuration mode, then the interface adopts the message digest key parameter specified by the area.

If this command is not specified in area configuration mode, then the interface adopts the message digest key parameter specified for the process.

If this command is not specified at any level, then OSPF MD5 authentication is disabled.

Command Modes

Interface configuration

Area configuration

XR Config mode

Virtual-link configuration

VRF configuration

Multi-area configuration

Sham-link configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

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  0/3/0/2
 message-digest-key 100 md5 OLD
```

You change the configuration to the following:

```
interface  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 HundredGigE 0/1/0/1
RP/0/RP0/CPU0:router(config-ospf-ar-if)# message-digest-key 19 md5 8ry4222
```

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

Command Default

LDP-IGP interface automatic configuration is disabled for OSPF.

Command Modes

Interface configuration

Area configuration

XR Config mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

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 o1
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**]

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

XR Config mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

Task ID

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 o1
RP/0/RP0/CPU0:router(config-ospf)# mpls ldp sync
```

mpls traffic-eng (OSPF)

To configure an Open Shortest Path First (OSPF) area for Multiprotocol Label Switching traffic engineering (MPLS TE), use the **mpls traffic-eng** command in the appropriate configuration mode. To remove the MPLS TE from an area, use the **no** form of this command.

mpls traffic-eng

Syntax Description

This command has no keywords or arguments.

Command Default

MPLS TE is not configured for OSPF.

Command Modes

Area configuration

Usage Guidelines

You must configure the **mpls traffic-eng** command for OSPF to support MPLS traffic engineering. OSPF provides the flooding mechanism that is used to flood TE link information.

Note This command is supported only in the default VRF mode.

We recommend that you configure the **mpls traffic-eng router-id** command instead of using the **router-id** command in mode.

OSPF support for MPLS TE is a component of the overall MPLS TE feature. Other MPLS TE software components must also be configured for this feature to be fully supported.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to associate loopback interface 0 with area 0, and area 0 is declared to be an MPLS area:

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router ospf 1
RP/0/RP0/CPU0:router(config-ospf)# router-id 10.10.10.10
RP/0/RP0/CPU0:router(config-ospf)# mpls traffic-eng router-id loopback 0
RP/0/RP0/CPU0:router(config-ospf)# area 0
RP/0/RP0/CPU0:router(config-ospf)# mpls traffic-eng
RP/0/RP0/CPU0:router(config-ospf-ar)# interface loopback 0
```

Related Commands

Command	Description
capability opaque disable	Controls the OSPF opaque LSA support capability.

Command	Description
mpls traffic-eng multicast intact (OSPF)	Specifies that the traffic engineering router identifier for the node is the IP address associated with a given interface.
router-id (OSPF)	Configures a router ID for the OSPF process.

mpls traffic-eng igp-intact (OSPF)

To ensure that the OSPF protocol installs at least one IPv4 next-hop when it adds the tunnel next-hops (igp-shortcuts), use the **mpls traffic-eng igp-intact** command in the router configuration mode. To disable IGP-intact, use the **no** form of this command.

mpls traffic-eng igp-intact

Command Default

IGP-intact is disabled.

Command Modes

Router configuration

Usage Guidelines

The OSPF protocol adds both tunnel next-hops (igp-shortcuts) and IPv4 next-hops to the next-hop list in the Routing Information Base (RIB), until it reaches the maximum number of paths. When IGP-intact is enabled, it ensures that the Routing Information Base (RIB) always has at least one IPv4 next-hop present in the list of next-hops when the number of paths is at maximum.

Note IGP-intact should be used only when Policy-Based Tunnel Selection (PBTS) is in use.

Configure IGP-intact only when Policy-Based Tunnel Selection (PBTS) is in use. This ensures that at least one IPv4 next-hop is available for the default Differentiated Services Code Point (DSCP) traffic class. It also prevents traffic loss for other DSCP traffic classes by diverting such traffic to an IPv4 next-hop when the corresponding tunnel is unavailable for forwarding.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to enable IGP-intact:

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router ospf 1
RP/0/RP0/CPU0:router(config-ospf)# mpls traffic-eng igp-intact
```

Related Commands

Command	Description
maximum paths (OSPF), on page 290	Configures the maximum number of parallel routes that the Open Shortest Path First (OSPF) protocol can support.

mpls traffic-eng multicast-intact (OSPF)

To enable multicast-intact for Open Shortest Path First (OSPF) routes so that multicast-intact paths are published to the Routing Information Base (RIB), use the **mpls traffic-eng multicast-intact** command in the appropriate configuration mode. To remove the MPLS TE area, use the **no** form of this command.

mpls traffic-eng multicast-intact

Syntax Description

This command has no keywords or arguments.

Command Default

MPLS TE is not configured for OSPF.

Command Modes

Router configuration

Usage Guidelines

OSPF support for MPLS TE is a component of the overall MPLS TE feature. Other MPLS TE software components must also be configured for this feature to be fully supported.

Note This command is supported only in the default VRF mode.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to enable publishing of multicast-intact paths to RIB:

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router ospf 1
RP/0/RP0/CPU0:router(config-ospf)# mpls traffic-eng multicast-intact
```

mpls traffic-eng ldp-sync-update (OSPF)

To enable Label Distribution Protocol (LDP)-OSPF Type 1 and Type 10 link-state advertisement (LSA) synchronization, use the **mpls traffic-eng ldp-sync-update** command in the appropriate mode. To disable LDP-LSA synchronization, use the **no** form of this command.

mpls traffic-eng ldp-sync-update

Command Default

LDP-LSA synchronization is disabled for OSPF.

Command Modes

Router configuration

Interface configuration

Area configuration

Usage Guidelines

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 o1
RP/0/RP0/CPU0:router(config-ospf)# mpls traffic-eng ldp-sync-update
```

mpls traffic-eng router-id (OSPF)

To specify that the traffic engineering router identifier for the node is the IP address associated with a given Open Shortest Path First (OSPF) interface, use the **mpls traffic-eng router-id** command in the appropriate configuration mode. To disable this feature, use the **no** form of this command.

mpls traffic-eng router-id { *router-id* | *type* *interface-path-id* }

Syntax Description

router-id	The 32-bit router ID value specified in four-part, dotted-decimal notation (must be in the valid IP address range of 0.0.0.0 to 255.255.255.255).
type	Interface type. For more information, use the question mark (?) online help function.
interface-path-id	Physical interface or virtual interface. Use the **show interfaces** command to see a list of all interfaces currently configured on the router. For more information about the syntax for the router, use the question mark (?) online help function.

Command Default

If this command is specified in router configuration mode, then the traffic engineering router identifier for the node is the IP address associated with a given interface.

Command Modes

Router configuration

Usage Guidelines

This identifier of the router acts as a stable IP address for the traffic engineering configuration. This IP address is flooded to all nodes. For all traffic engineering tunnels originating at other nodes and ending at this node, you must set the tunnel destination to the traffic engineering router identifier of the destination node, because that is the address that the traffic engineering topology database at the tunnel head uses for its path calculation.

Note We recommend that loopback interfaces be used for Multiprotocol Label Switching traffic engineering (MPLS TE), because they are more stable than physical interfaces.

Note This command is supported only in the default VRF mode.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to specify the traffic engineering router identifier as the IP address associated with loopback interface 0:

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router ospf 1
RP/0/RP0/CPU0:router(config-ospf)# mpls traffic-eng router-id loopback 0
```

Related Commands

Command	Description
mpls traffic-eng (OSPF), on page 298	Configures an OSPF area for MPLS TE.

mtu-ignore (OSPF)

To prevent Open Shortest Path First (OSPF) from checking whether neighbors are using the same maximum transmission unit (MTU) on a common interface when exchanging database descriptor (DBD) packets, use the **mtu-ignore** command in the appropriate mode. To reset to default, use the **no** form of this command.

mtu-ignore [**disable** | **enable**]

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

Usage Guidelines

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*

Syntax Description

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

Command Default

An OSPF network is enabled for one area only.

Command Modes

Area configuration

Command History

Release	Modification
Release 24.2.11	The command is extended to support loopback interfaces.
Release 6.0	This command was introduced.

Usage Guidelines

Use the **multi-area-interface** command to enable area border routers (ABRs) to establish multiple adjacencies for different OSPF areas.

Each multiple area adjacency is announced as a point-to-point unnumbered link in the configured area. This point-to-point link provides a topological path for that area. The first or primary adjacency using the link advertises the link consistent with draft-ietf-ospf-multi-area-adj-06.txt.

You can configure multi-area adjacency on any interface where only two OSF speakers are attached. In the case of native broadcast networks, the interface must be configured as an OPSF point-to-point type using the **network point-to-point** command to enable the interface for a multi-area adjacency.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to enable multiple area adjacency for OSPF 109:

```
RP/0/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)#
```

The following example shows how to configure, associate, and reuse a loopback IPv4 interface address for all areas in an ABR:

```
Router(config)#interface Loopback 0
Router(config-if)#ipv4 address 192.0.2.2/32
Router(config-if)#commit

Router(config)#router ospf 100
Router(config-ospf)#router-id 10.10.10.1
Router(config-ospf)#area 1
Router(config-ospf-ar)#interface Loopback 0
Router(config-ospf-ar-if)#commit

Router(config)#router ospf 100
Router(config-ospf)#router-id 10.10.10.1
Router(config-ospf)#area 0
Router(config-ospf-ar)#multi-area-interface Loopback 0
Router(config-ospf-ar-mif)#commit
```

Related Commands

Command	Description
show ospf interface, on page 374	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*]

Syntax Description

ip-address	Interface IP address of the neighbor.
cost *number*	(Optional) Assigns a cost to the neighbor, in the form of an integer from 1 to 65535. Neighbors with no specific cost configured assume the cost of the interface, based on the **cost** command. On point-to-multipoint interfaces, **cost** *number* is the only keyword and argument combination that works. The **cost** keyword does not apply to nonbroadcast multiaccess (NBMA) networks.
priority *number*	(Optional) Specifies an 8-bit number indicating the router priority value of the nonbroadcast neighbor associated with the IP address specified. The **priority** keyword does not apply to point-to-multipoint interfaces.
poll-interval *seconds*	(Optional) Specifies an unsigned integer value (in seconds) reflecting the poll interval. RFC 1247 recommends that this value be much larger than the hello interval. The **poll-interval** keyword does not apply to point-to-multipoint interfaces.

Command Default

No configuration is specified.

priority *number* : 0

poll-interval *seconds* : 120 seconds (2 minutes)

Command Modes

Interface configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

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 HundredGigE 1/0/0/3
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 HundredGigE 1/0/0/3
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
```

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**

Syntax Description

ip-address	IP address of the neighbor to which outgoing LSAs are blocked.

Command Default

Instead of all outgoing LSAs being filtered to the neighbor, they are flooded to the neighbor.

Command Modes

Interface configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

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 HundredGigE 1/0/0/3
RP/0/RP0/CPU0:router(config-ospf-ar-if)# neighbor 10.2.3.4 database-filter all out
```

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** } }

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.

TenGigEthernet interfaces default to broadcast.

Command Modes

Interface configuration

Area configuration

Router configuration

VRF configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

Use the **network** command to configure broadcast networks as NBMA networks when, for example, routers in your network do not support multicast addressing.

Configuring NBMA networks as either broadcast or nonbroadcast assumes that there are virtual circuits from every router to every router or fully meshed network. However, there are other configurations where this assumption is not true; for example, a partially meshed network. In these cases, you can configure the OSPF network type as a point-to-multipoint network. Routing between two routers that are not directly connected go through the router that has virtual circuits to both routers. You need not configure neighbors when using this command.

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/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 HundredGigE 0/1/0/3
RP/0/RP0/CPU0:router(config-ospf-ar-if)# network non-broadcast
RP/0/RP0/CPU0:router(config-ospf-ar-if)# neighbor 172.16.3.4 priority 1 poll-interval 180
```

nsf (OSPF)

To configure nonstop forwarding (NSF) for the Open Shortest Path First (OSPF) protocol, use the **nsf** command in the appropriate mode. To remove this command from the configuration file and restore the system to its default condition, use the **no** form of this command.

nsf { **cisco** [**enforce global**] | **ietf** [**helper disable**] }

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

XR Config mode

VRF configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

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:

```
Router# configure
Router(config)# router ospf 1
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*

Syntax Description

seconds	Length of time (in seconds) allowed for NSF external route queries. Range is 1 to 3600 seconds.

Command Default

seconds : 300

Command Modes

Router configuration

VRF configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

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*

Syntax Description

seconds	Length of time (in seconds) between consecutive restart attempts. Range is 90 to 3600 seconds.

Command Default

seconds : 90

Command Modes

Router configuration

VRF configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

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:routerr(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*

Syntax Description

seconds	The length of time (in seconds) that routes are held in the RIB. Range is 90 to 3600 seconds.

Command Default

seconds : 95

Command Modes

Router configuration

VRF configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

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
```

nsr (OSPF)

To configure nonstop routing (NSR) for the Open Shortest Path First (OSPF) protocol, use the **nsr** command in OSPF router configuration mode. To remove this command from the configuration file, use the **no** form of this command.

nsr [**disable**]
no nsr [**disable**]

Syntax Description

disable	(Optional) Disables NSR for all VRFs in this process.

Command Default

NSR is enabled.

Command Modes

Router configuration

Usage Guidelines

The NSR feature allows an OSPF process on the active RP to synchronize all necessary data and states with the OSPF process on the standby RP. When the switchover happens, the OSPF process on the newly active RP has all the necessary data and states to continue running and does not require any help from its neighbors.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to configure NSR:

```
Router# configure
Router(config)# router ospf 1
Router(config-ospf)# nsr
```

The following example shows how to disable NSR:

```
Router# configure
Router(config)# router ospf 1
Router(config-ospf)# nsr disable
```

nssa (OSPF)

To configure an area as a not-so-stubby area (NSSA), use the **nssa** command in area configuration mode. To remove the NSSA distinction from the area, use the **no** form of this command.

nssa [**no-redistribution**] [**default-information-originate** [**metric** *metric-value*] [**metric-type** *type-value*]] [**no-summary**]
no nssa

Syntax Description

no-redistribution	(Optional) Imports routes only into the normal areas, but not into the NSSA area, by the **redistribute** command when the router is an NSSA Area Border Router (ABR).
default-information-originate	(Optional) Generates a Type 7 default into the NSSA area. This keyword takes effect only on an NSSA ABR or NSSA Autonomous System Boundary Router (ASBR).
metric *metric-value*	(Optional) Specifies the metric used for generating the default route. If you omit a value and do not specify a value using the **defaultmetric** command, the default metric value is 10. Range is 1 to 16777214.
metric-type *type-value*	(Optional) Specifies the external link type associated with the default route advertised into the OSPF routing domain. It can be one of the following values: **1**—Type 1 external route **2**—Type 2 external route
no-summary	(Optional) Prevents an ABR from sending summary link advertisements into the NSSA.

Command Default

No NSSA area is defined.

Command Modes

Area configuration

Usage Guidelines

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.

Examples

The following example shows how to configure area 1 as an NSSA area:

```
Router# configure
Router(config)# router ospf 1
Router(config-ospf)# area 1
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

Command Default

Routers are displayed by router ID or neighbor ID.

Command Modes

XR Config mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

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*

Syntax Description

bytes	Size, in bytes. Range is 576 to 10000 bytes.

Command Default

If the command is not specified, the default packet size is either the interface IP MTU size (if that is lower than 9000 bytes) or 9000 bytes.

Command Modes

XR Config mode

Area configuration

Interface configuration

VRF configuration

Multi-area configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

Use the **packet-size** command to customize the size of OSPF packets. The OSPF protocol compares the packet size and the MTU size and uses the lower packet size value.

If the command is not configured, the default packet size is equal to the interface IP MTU size (if that is lower than 9000 bytes) or 9000 bytes. For example, if the interface IP MTU size is 1500 bytes, OSPF uses packet size of 1500 bytes on the interface because the byte size is lower than 9000 bytes. If the interface IP MTU size is 9500 bytes, OSPF uses packet size of 9000 bytes on the interface because the byte size exceeds 9000 bytes. The interface IP MTU size depends on the interface and the platform. In most cases, the default interface IP MTU value will be lower than 9000 bytes.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to configure the packet size on an interface:

```
RP/0/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 HundredGigE 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**]

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

XR Config mode

VRF configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

OSPF routing information is neither sent nor received through the specified interface. The interface appears as a stub network in the OSPF router (Type 1) link-state advertisement (LSA).

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows that GigabitEthernet interface 1/0/0/2 reduces OSPF updates because passive mode is enabled; however, HundredGigE 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 HundredGigE 1/0/0/2
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 HundredGigE 1/0/0/3
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*

Syntax Description

value	8-bit unsigned integer indicating the router priority value. Range is 0 to 255.

Command Default

If this command is not specified in interface configuration mode, then the interface adopts the priority parameter specified by the area.

If this command is not specified in area configuration mode, then the interface adopts the priority parameter specified for the process.

If this command is not specified at any level, then the default priority is 1.

Command Modes

Interface configuration

Area configuration

Router configuration

VRF configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

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:

```
Router A
  RP/0/RP0/CPU0:router# configure
  RP/0/RP0/CPU0:router(config)# interface HundredGigE 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 HundredGigE 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

Router B
  RP/0/RP0/CPU0:router# configure
  RP/0/RP0/CPU0:router(config)# interface HundredGigE 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 HundredGigE 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
```

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 XR Config mode. To reenable the OSPF protocol, use the **no** form of this command.

protocol shutdown

Command Default

No default behavior or values

Command Modes

XR Config mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

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 XR Config mode. To return to the system default value, use the **no** form of this command.

queue dispatch flush-lsa *count*

Syntax Description

count	Maximum number of LSAs flushed per run. Range is 30 to 3000.

Command Default

The default LSAs flushed per run is 150 (when the count is not configured).

Command Modes

XR Config mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

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 command to see the queue dispatch values, peak lengths, and limits.

queue dispatch incoming

To limit the number of incoming packets (LSAUpdates, LSAcks, DBDs, LSRequests, and Hellos that trigger a change state) processed, use the **queue dispatch incoming** command in XR Config mode. To return to the system default value, use the **no** form of this command.

queue dispatch incoming *count*

Syntax Description

count	Maximum number of continuous events processed. Range is 30 to 3000.

Command Default

The default incoming count is 300 packets (when the count is not configured).

Command Modes

XR Config mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

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 381 command to see 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 XR Config mode. To return to the system default value, use the **no** form of this command.

queue dispatch rate-limited-lsa *count*

Syntax Description

count	Maximum number of rate-limited LSAs processed per run. Range is 30 to 3000.

Command Default

The default number of rate-limited LSAs processed per run is 300 (when this count is not configured).

Command Modes

XR Config mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to set the maximum number of rate-limited LSA (re-)originations processed per run to 300:

```
RP/0/RP0/CPU0:router(config-ospf)# queue dispatch rate-limited-lsa 300
```

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 XR Config mode. To return to the system default value, use the **no** form of this command.

queue dispatch spf-lsa-limit *count*

Syntax Description

count	Maximum number of continuous Type 3-4 and Type 5-7 LSAs processed per SPF in each scheduled iteration within a single SPF run. Range is 30 to 3000.

Command Default

The default number of Type 3-4 and Type 5-7 processed per run is 150 LSAs (when this command is not configured).

Command Modes

XR Config mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

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
```

queue limit

To set the high watermark for incoming events by priority, use the **queue limit** in XR Config mode. To return to the system default values, use the **no** form of this command.

queue limit { **high** | **medium** | **low** } *count*

Syntax Description

high	High watermark for incoming high-priority events (state-changing Hellos).
medium	High watermark for incoming medium-priority events (LSA ACK).
low	High watermark for incoming low-priority events (DBD/LSUpd/LSReq).
count	Maximum number of events per queue. Events are dropped when the priority queue size exceeds this value. Range is 1000 to 30000.

Command Default

High watermark: 9500 (when the corresponding configuration is not present).

Medium watermark: 9000 (when the corresponding configuration is not present).

Low watermark: 8000 (when the corresponding configuration is not present).

Command Modes

XR Config mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

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
```

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**]

Syntax Description

ip-address	IP address in four-part, dotted-decimal notation.
mask	IP address mask.
advertise	(Optional) Sets the address range status to advertise and generates a Type 3 summary link-state advertisement (LSA).
not-advertise	(Optional) Sets the address range status to DoNotAdvertise. The Type 3 summary LSA is suppressed and the component networks remain hidden from other networks.

Command Default

When this command is not specified for Area Border Routers (ABRs), routes at an area boundary are not consolidated or summarized.

Advertise is the default.

Command Modes

Area configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

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 HundredGigE 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 HundredGigE 0/1/0/0
RP/0/RP0/CPU0:router(config-ospf-ar-if)# interface HundredGigE 0/1/0/0
RP/0/RP0/CPU0:router(config-ospf-ar-if)# interface HundredGigE 0/1/0/1
RP/0/RP0/CPU0:router(config-ospf-ar-if)# interface HundredGigE 0/1/0/2
RP/0/RP0/CPU0:router(config-ospf-ar-if)# interface HundredGigE 0/1/0/3
RP/0/RP0/CPU0:router(config-ospf-ar-if)# interface HundredGigE 0/2/0/0
RP/0/RP0/CPU0:router(config-ospf-ar-if)# interface HundredGigE 0/2/0/1
RP/0/RP0/CPU0:router(config-ospf-ar-if)# interface HundredGigE 0/2/0/2
RP/0/RP0/CPU0:router(config-ospf-ar-if)# interface HundredGigE 0/2/0/3
RP/0/RP0/CPU0:router(config-ospf-ar-if)# end
```

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*]

Local Interface Routes
redistribute connected [**instance** *instance-name*] [**instance IPCP**] [**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*]

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*]

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*]

Routing Information Protocol (RIP)
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*]

Syntax Description

bgp	Distributes routes from the BGP protocol.

process-id	For the **bgp** keyword, an autonomous system number has the following ranges: • Range for 2-byte Autonomous system numbers (ASNs) is 1 to 65535. • Range for 4-byte Autonomous system numbers (ASNs) in asplain format is 1 to 4294967295. • Range for 4-byte Autonomous system numbers (ASNs) is asdot format is 1.0 to 65535.65535. For the **isis** keyword, an IS-IS instance name from which routes are to be redistributed. The value takes the form of a string. A decimal number can be entered, but it is stored internally as a string. For the **ospf** keyword, an OSPF instance name from which routes are to be redistributed. The value takes the form of a string. A decimal number can be entered, but it is stored internally as a string.
preserve-med	(Optional) Preserves the Multi Exit Discriminator (MED) of BGP routes.
metric *metric-value*	(Optional) Specifies the metric used for the redistributed route. Range is 1 to 16777214. Use a value consistent with the source protocol.
metric-type { **1** | **2** }	(Optional) Specifies the external link type associated with the route advertised into the OSPF routing domain. It can be one of two values: • **1** —Type 1 external route • **2** —Type 2 external route
tag *tag-value*	(Optional) Specifies the value attached to each external route. This value is not used by the OSPF protocol itself, but is carried in the external LSAs. Range is 0 to 4294967295.
route-policy *policy-name*	(Optional) Specifies the identifier of a configured policy. A policy is used to filter the importation of routes from this source routing protocol to OSPF.
connected	Distributes routes that are established automatically by virtue of having enabled IP on an interface.
instance	Connected instance.
instance-name	Name of the connected instance.
instance IPCP	Distributes routes from IPCP protocols.
isis	Distributes routes from the IS-IS protocol.
level-1	(Optional) Redistributes Level 1 routes into other IP routing protocols independently.
l **evel-1-2**	(Optional) Distributes both Level 1 and Level 2 routes into other IP routing protocols.
level-2	(Optional) Distributes Level 2 routes into other IP routing protocols independently.

ospf	Distributes routes from the OSPF protocol.
match { **internal** | **external [1** | **2**] | **nssa-external** [**1** | **2]** }	(Optional) Specifies the criteria by which OSPF routes are redistributed into other routing domains. It can be one or more of the following: • **internal**—Routes that are internal to a specific autonomous system (intra- and inter-area OSPF routes). • **external** [**1 | 2**]—Routes that are external to the autonomous system, but are imported into OSPF as Type 1 or Type 2 external routes. • **nssa-external** [**1 | 2**]—Routes that are external to the autonomous system, but are imported into OSPF as Type 1 or Type 2 not-so-stubby area (NSSA) external routes. For the **external** and **nssa-external** options, if a type is not specified, then both Type 1 and Type 2 are assumed. If no match is specified, the default is no filtering.
rip	Distributes routes from the RIP protocol.
static	Distributes IP static routes.
dagr	Distributes routes from the directed-attached gateway redundancy (DAGR).

Command Default

Route redistribution is disabled.

metric *metric-value:* Default is 20 for routes from all protocols except BGP routes, for which the default is 1.

metric-type : Type 2 external route.

Command Modes

Router configuration

VRF configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

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 Router Routing Command Reference Guide*.

Whenever you use the **redistribute** or default-information originate (OSPF) 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 HundredGigE 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 HundredGigE 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 HundredGigE 0/2/0/2
!
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 HundredGigE 0/1/0/1
```

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*

Syntax Description

seconds	Time (in seconds) between retransmissions. It must be greater than the expected round-trip delay between any two routers on the attached network. Range is 1 to 65535 seconds.

Command Default

If this command is not specified in interface configuration mode, then the interface adopts the retransmit interval parameter specified by the area.

If this command is not specified in area configuration mode, then the interface adopts the retransmit interval parameter specified for the process.

If this command is not specified at any level, then the default retransmit interval is 5 seconds.

Command Modes

Interface configuration

Area configuration

XR Config mode

Virtual-link configuration

VRF configuration

Multi-area configuration

Sham-link configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

When a router sends an LSA to its neighbor, it keeps the LSA until it receives the acknowledgment message. If the router receives no acknowledgment, it resends the LSA.

The setting of this parameter should be conservative, or needless retransmission results. The value should be larger for serial lines and virtual links.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to set the retransmit interval value to 8 seconds in interface configuration mode:

```
RP/0/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 HundredGigE 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** }

Syntax Description

route-policy-name	Name of route policy.
in	Applies policy to inbound routes.
out	Applies policy to outbound routes.

Command Default

No policy is applied.

Command Modes

Area configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

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

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to specify an OSPF route policy for inbound routes in area 0:

```
RP/0/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*

Syntax Description

router-id	32-bit router ID value specified in four-part, dotted-decimal notation.

Command Default

If this command is not configured, the router ID is the highest IP version 4 (IPv4) address for an interface on the router, with any loopback interface taking precedence.

Command Modes

Router configuration

VRF configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

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.

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
```

router ospf (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*

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 mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

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.

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
```

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*]

Syntax Description

hops *hops-number*	IP hops. Maximum number of hops allowed. Range is 1 to 254 hops.

Command Default

hops-number : 1

Command Modes

XR Config mode

Area configuration

Interface configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

The **security ttl** command is used for the Generalized TTL Security Mechanism (GTSM) feature to prevent network attacks.

During the act of receiving Link State Advertisement (LSA) from neighbors, network attacks can occur because there are no checks that unicast or multicast packets are originating from a neighbor that is one hop away or multiple hops away over virtual links.

For virtual links, OSPF packets travel multiple hops across the network; hence, the TTL value can be decremented several times. For these type of links, a minimum TTL value must be allowed and accepted for multiple-hop packets.

To filter network attacks originating from invalid sources traveling over multiple hops, the GTSM, RFC 3682, is used to prevent the attacks. GTSM filters link-local addresses and allows for only one-hop neighbor adjacencies through the configuration of TTL value 255. The TTL value in the IP header is set to when OSPF packets are originated and checked on the received OSPF packets against the default GTSM TTL value 255 or the user configured GTSM TTL value, blocking unauthorized OSPF packets originated from TTL hops away.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to set the security TTL for an interface:

```
RP/0/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 HundredGigE 0/6/0/3
RP/0/RP0/CPU0:router(config-ospf-ar-if)# security ttl 2
```

segment-routing sr-prefer prefix-list

To set the preference of segment routing (SR) labels over label distribution protocol (LDP) labels, use the **segment-routing sr-prefer prefix-list** command. The label preference is used to decide the labeled path that will be used in case both LDP and SR labeled paths are available. This only affects the traffic that comes as IP (un-labeled) or traffic that resolves over a labeled path (such as L3VPNs, P2VPNs and so on).

Note If ACL is used, OSPF signals preference of SR labels over LDP labels only for prefixes that match the ACL. If ACL is not used, SR labels preference is signaled for all prefixes.

segment-routing sr-prefer prefix-list *[acl-name]*

Syntax Description

prefix-list	Sets preference of SR labels over LDP labels.
[acl-name]	Name of access control list.

Command Default

LDP labels are preferred by default.

Command Modes

Router configuration

Area configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced

Usage Guidelines

No specific guidelines impact the use of this command.

Task ID

Task ID	Operations
ospf	read, write

Examples

```
RP/0/RP0/CPU0:router#configure
RP/0/RP0/CPU0:router(config)#ipv4 prefix-list foo
RP/0/RP0/CPU0:router(config)#router ospf 1
RP/0/RP0/CPU0:router(config-ospf)# segment-routing sr-prefer prefix-list foo
  area 0
   interface Loopback0
    prefix-sid index 1
   !
   interface GigabitEthernet0/0/0/0
   !
   interface GigabitEthernet0/2/0/0
```

```
 !
 interface GigabitEthernet0/2/0/3
 !
!
area 1
 interface GigabitEthernet0/2/0/7
 !
```

sham-link

To configure an Open Shortest Path First OSPF sham link between two provider edge routers, use the **sham-link** command in VRF area configuration mode. To terminate an (OSPF) sham link, use the **no** form of this command.

sham-link *source-address* *destination-address*

Syntax Description

source-address	IP address of the local (source) sham-link endpoint specified in four-part, dotted-decimal notation.
destination-address	IP address of the remote (destination) sham-link endpoint specified in four-part, dotted-decimal notation.

Command Default

No sham link is configured.

Command Modes

VRF area configuration.

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

Use the **sham-link** command to configure a point-to-point connection between two provider edge (PE) routers creating an interconnect between two VPN sites (VPN backbone). Sham links are configured on PE provider edge (PE) routers in a Multiprotocol Label Switching (MPLS) VPN backbone.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to configure an OSPF sham link:

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router ospf 109
RP/0/RP0/CPU0:router(config_ospf)# vrf vrf_a
RP/0/RP0/CPU0:router(config_ospf_vrf)# area 0
RP/0/RP0/CPU0:router(config_ospf_vrf_ar)# sham-link 192.168.40.0 172.16.30.0
RP/0/RP0/CPU0:router(config_ospf_vrf_ar_sl)# cost 23
```

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*] [**vrf** {*vrf-name* | **all**}] [**summary**]

Syntax Description

process-name	(Optional) Name that uniquely identifies an OSPF routing process. The process name is defined by the **router ospf** command. If this argument is included, only information for the specified routing process is displayed.
vrf *vrf-name* **all**	(Optional) Specifies an OSPF VPN routing and forwarding (VRF) instance. The *vrf-name* argument can be specified as an arbitrary string. The strings "default" and "all" are reserved values of the *vrf-name* argument.
summary	(Optional) Displays OSPF summary information.

Command Default

IPv4 and unicast address prefixes

Command Modes

XR EXEC mode

OSPFv3

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

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 msecs
 Minimum hold time between two consecutive SPFs 10000 msecs
 Maximum wait time between two consecutive SPFs 10000 msecs
 Initial LSA throttle delay 500 msecs
 Minimum hold time for LSA throttle 5000 msecs
 Maximum wait time for LSA throttle 5000 msecs
 Minimum LSA interval 5000 msecs. Minimum LSA arrival 1 secs
```

```
Maximum number of configured interfaces 255
Number of external LSA 0. Checksum Sum 00000000
Number of opaque AS LSA 0. Checksum Sum 00000000
Number of DCbitless external and opaque AS LSA 0
Number of DoNotAge external and opaque AS LSA 0
Number of areas in this router is 2. 2 normal 0 stub 0 nssa
External flood list length 0
Non-Stop Forwarding enabled
   Area BACKBONE(0) (Inactive)
       Number of interfaces in this area is 2
       SPF algorithm executed 8 times
       Number of LSA 2. Checksum Sum 0x01ba83
       Number of opaque link LSA 0. Checksum Sum 00000000
       Number of DCbitless LSA 0
       Number of indication LSA 0
       Number of DoNotAge LSA 0
       Flood list length 0
   Area 1
       Number of interfaces in this area is 1
       SPF algorithm executed 9 times
       Number of LSA 2. Checksum Sum 0x0153ea
       Number of opaque link LSA 0. Checksum Sum 00000000
       Number of DCbitless LSA 0
       Number of indication LSA 0
       Number of DoNotAge LSA 0
       Flood list length 0
```

This table describes the significant fields shown in the display.

Table 41: 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.

Field	Description
Number of areas	Number of areas in router, area addresses, and so on.
Area BACKBONE	Backbone is area 0.

This sample output from the **show ospf vrf** *vrf_name* command displays the VRF Lite status:

```
RP/0/RP0/CPU0:router#show ospf vrf vrf1

 VRF vrf1 in Routing Process "ospf 100" with ID 1.1.1.1
 NSR (Non-stop routing) is Disabled
 Supports only single TOS(TOS0) routes
 Supports opaque LSA
 It is an area border router
 VRF Lite is enabled
 Router is not originating router-LSAs with maximum metric
 Initial SPF schedule delay 50 msecs
 Minimum hold time between two consecutive SPFs 200 msecs
 Maximum wait time between two consecutive SPFs 5000 msecs
 Initial LSA throttle delay 50 msecs
 Minimum hold time for LSA throttle 200 msecs
 Maximum wait time for LSA throttle 5000 msecs
 Minimum LSA interval 200 msecs. Minimum LSA arrival 100 msecs
 LSA refresh interval 1800 seconds
 Flood pacing interval 33 msecs. Retransmission pacing interval 66 msecs
 Adjacency stagger enabled; initial (per area): 2, maximum: 64
    Number of neighbors forming: 0, 2 full
 Maximum number of configured interfaces 1024
 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
 SNMP trap is disabled
    Area BACKBONE(0)
        Number of interfaces in this area is 1
        SPF algorithm executed 4 times
        Number of LSA 16.  Checksum Sum 0x071c6a
        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
        Number of LFA enabled interfaces 0, LFA revision 0
        Number of Per Prefix LFA enabled interfaces 0
        Number of neighbors forming in staggered mode 0, 1 full
    Area 1
        Number of interfaces in this area is 4
        SPF algorithm executed 5 times
        Number of LSA 14.  Checksum Sum 0x066d93
        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
        Number of LFA enabled interfaces 0, LFA revision 0
        Number of Per Prefix LFA enabled interfaces 0
        Number of neighbors forming in staggered mode 0, 1 full
```

The **show ospf** command displays the maximum number of redistributed prefix is limited to 1000.

```
Router #show ospf
Thu Dec  8 18:16:48.332 IST

 Routing Process "ospf 1" with ID 192.168.0.1
 Role: Primary Active
 NSR (Non-stop routing) is Enabled
 Supports only single TOS(TOS0) routes
 Supports opaque LSA
 It is an autonomous system boundary router
 Maximum number of non self-generated LSA allowed 1000
    Current number of non self-generated LSA 804
    Threshold for warning message 60%
    Ignore-time 1 minutes, reset-time 2 minutes
    Ignore-count allowed 2, current ignore-count 0
 Redistributing External Routes from,
    static
 Maximum number of redistributed prefixes 1000
    Threshold for warning message 70%
    Current number of redistributed prefixes 100
```

show ospf border-routers

To display the internal Open Shortest Path First (OSPF) routing table entries to an Area Border Router (ABR) and Autonomous System Boundary Router (ASBR), use the **show ospf border-routers** command in XR EXEC mode.

show ospf [*process-name*] [**vrf** {*vrf-name* | **all**}] **border-routers** [*router-id*]

Syntax Description

process-name	(Optional) OSPF process name. If this argument is included, only information for the specified routing process is included.
vrf *vrf-name* **all**	(Optional) Specifies an OSPF VPN routing and forwarding (VRF) instance. The *vrf-name* argument can be specified as an arbitrary string. The strings "default" and "all" are reserved vrf-names.
router-id	(Optional) Router ID associated with the border router. The value of the *router-id* argument can be any 32-bit router ID value specified in four-part, dotted-decimal notation. No default exists.

Command Default

IPv4 and unicast address prefixes

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

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 42: show ospf border-routers Field Descriptions

Field	Description
i	Type of this route; i indicates an intra-area route, I an interarea route.
172.31.97.53	Router ID of destination.
[1]	Cost of using this route.
172.16.1.53	Next-Next hop toward the destination.
GigabitEthernet 3/0/0/0	Packets destined for 172.16.1.53 are sent over GigabitEthernet interface 3/0/0/0.
ABR/ASBR	Router type of the destination; it is either an Area Border Router (ABR) or Autonomous System Boundary Router (ASBR) or both.
Area 0	Area ID of the area from which this route is learned.
SPF 3	Internal number of the shortest path first (SPF) calculation that installs this route.

show ospf database

To display lists of information related to the Open Shortest Path First (OSPF) database for a specific router, use the **show ospf database** command in XR EXEC mode.

show ospf [*process-name*] [**vrf** {*vrf-name* | **all**}] [*area-id*] **database**
show ospf [*process-name*] [**vrf** {*vrf-name* | **all**}] [*area-id*] **database** [**adv-router** *ip-address*]
show ospf [*process-name*] [**vrf** {*vrf-name* | **all**}] [*area-id*] **database** [**asbr-summary**] [*link-state-id*]
show ospf [*process-name*] [**vrf** {*vrf-name* | **all**}] [*area-id*] **database** [**asbr-summary**] [*link-state-id*] [**internal**] [**adv-router** [*ip-address*]]
show ospf [*process-name*] [**vrf** {*vrf-name* | **all**}] [*area-id*] **database** [**asbr-summary**] [*link-state-id*] [**internal**] [**self-originate**]
show ospf [*process-name*] [**vrf** {*vrf-name* | **all**}] [*area-id*] **database** [**database-summary**]
show ospf [*process-name*] [**vrf** {*vrf-name* | **all**}] [*area-id*] **database** [**external**] [*link-state-id*]
show ospf [*process-name*] [**vrf** {*vrf-name* | **all**}] [*area-id*] **database** [**external**] [*link-state-id*] [**internal**] [**adv-router** [*ip-address*]]
show ospf [*process-name*] [**vrf** {*vrf-name* | **all**}] [*area-id*] **database** [**external**] [*link-state-id*] [**internal**] [**self-originate**]
show ospf [*process-name*] [**vrf** {*vrf-name* | **all**}] [*area-id*] **database** [**network**] [*link-state-id*]
show ospf [*process-name*] [**vrf** {*vrf-name* | **all**}] [*area-id*] **database** [**network**] [*link-state-id*] [**internal**] [**adv-router** [*ip-address*]]
show ospf [*process-name*] [**vrf** {*vrf-name* | **all**}] [*area-id*] **database** [**network**] [*link-state-id*] [**internal**] [**self-originate**]
show ospf [*process-name*] [**vrf** {*vrf-name* | **all**}] [*area-id*] **database** [**nssa-external**] [*link-state-id*]
show ospf [*process-name*] [**vrf** {*vrf-name* | **all**}] [*area-id*] **database** [**nssa-external**] [*link-state-id*] [**internal**] [**adv-router** [*ip-address*]]
show ospf [*process-name*] [**vrf** {*vrf-name* | **all**}] [*area-id*] **database** [**nssa-external**] [*link-state-id*] [**internal**] [**self-originate**]
show ospf [*process-name*] [**vrf** {*vrf-name* | **all**}] [*area-id*] **database** [**opaque-area**] [*link-state-id*]
show ospf [*process-name*] [**vrf** {*vrf-name* | **all**}] [*area-id*] **database** [**opaque-area**] [*link-state-id*] [**internal**] [**adv-router**] [*ip-address*]
show ospf [*process-name*] [**vrf** {*vrf-name* | **all**}] [*area-id*] **database** [**opaque-area**] [*link-state-id*] [**internal**] [**self-originate**]
show ospf [*process-name*] [**vrf** {*vrf-name* | **all**}] [*area-id*] **database** [**opaque-as**] [*link-state-id*]
show ospf [*process-name*] [**vrf** {*vrf-name* | **all**}] [*area-id*] **database** [**opaque-as**] [*link-state-id*] [**internal**] [**adv-router** [*ip-address*]]
show ospf [*process-name*] [**vrf** {*vrf-name* | **all**}] [*area-id*] **database** [**opaque-as**] [*link-state-id*] [**internal**] [**self-originate**]
show ospf [*process-name*] [**vrf** {*vrf-name* | **all**}] [*area-id*] **database** [**opaque-link**] [*link-state-id*]
show ospf [*process-name*] [**vrf** {*vrf-name* | **all**}] [*area-id*] **database** [**opaque-link**] [*link-state-id*] [**internal**] [**adv-router** [*ip-address*]]
show ospf [*process-name*] [**vrf** {*vrf-name* | **all**}] [*area-id*] **database** [**opaque-link**] [*link-state-id*] [**internal**] [**self-originate**]
show ospf [*process-name*] [**vrf** {*vrf-name* | **all**}] [*area-id*] **database** [**router**] [*link-state-id*]
show ospf [*process-name*] [**vrf** {*vrf-name* | **all**}] [*area-id*] **database** [**router**] [**internal**] [**adv-router** [*ip-address*]]
show ospf [*process-name*] [**vrf** {*vrf-name* | **all**}] [*area-id*] **database** [**router**] [**internal**] [**self-originate**] [*link-state-id*]
show ospf [*process-name*] [**vrf** {*vrf-name* | **all**}] [*area-id*] **database** [**self-originate**]

show ospf [*process-name*] [**vrf** {*vrf-name* | **all**}] [*area-id*] **database** [**summary**] [*link-state-id*]
show ospf [*process-name*] [**vrf** {*vrf-name* | **all**}] [*area-id*] **database** [**summary**] [*link-state-id*] [**internal**] [**adv-router** [*ip-address*]]
show ospf [*process-name*] [**vrf** {*vrf-name* | **all**}] [*area-id*] **database** [**summary**] [*link-state-id*] [**internal**] [**self-originate**] [*link-state-id*]

Syntax Description

process-name	(Optional) OSPF process name that uniquely identifies an OSPF routing process. The process name is any alphanumeric string no longer than 40 characters. If this argument is included, only information for the specified routing process is included.
vrf	(Optional) Specifies an OSPF VPN routing and forwarding (VRF) instance.
vrf-name	(Optional) Name of the OSPF VRF. The *vrf-name* argument can be specified as an arbitrary string. The strings “default” and “all” are reserved VRF names.
all	(Optional) Specifies all OSPF VRF instances.
area-id	(Optional) Area number used to define the particular area.
adv-router *ip-address*	(Optional) Displays all LSAs of the specified router.
asbr-summary	(Optional) Displays information only about the Autonomous System Boundary Router (ASBR) summary LSAs.
link-state-id	(Optional) Portion of the Internet environment that is being described by the advertisement. The value entered depends on the link-state type of the advertisement. It must be entered in the form of an IP address. When the link-state advertisement (LSA) is describing a network, the *link-state-id* can take one of two forms: • The network IP address (as in Type 3 summary link advertisements and in autonomous system external link advertisements). • A derived address obtained from the link-state ID. **Note** Masking the link-state ID of a network link advertisement with the subnet mask of the network yields the IP address of the network. When the LSA is describing a router, the link-state ID is always the OSPF router ID of the described router. When an autonomous system external advertisement (LS Type = 5) is describing a default route, its link-state ID is set to Default Destination (0.0.0.0).
internal	(Optional) Displays internal LSA information.
self-originate	(Optional) Displays only self-originated LSAs (from the local router).
database-summary	(Optional) Displays how many of each type of LSA for each area there are in the database and the total.
external	(Optional) Displays information only about the external LSAs.
network	(Optional) Displays information only about the network LSAs.

nssa-external	(Optional) Displays information only about the not-so-stubby area (NSSA) external LSAs.
opaque-area	(Optional) Displays information about the opaque Type 10 LSAs. Type 10 denotes an area-local scope. Refer to RFC 2370 for more information on the opaque LSA options.
opaque-as	(Optional) Displays information about the opaque Type 11 LSAs. Type 11 denotes that the LSA is flooded throughout the autonomous system.
opaque-link	(Optional) Displays information about the opaque Type 9 LSAs. Type 9 denotes a link-local scope.
router	(Optional) Displays information only about the router LSAs.
summary	(Optional) Displays information only about the summary LSAs.
detail	(Optional) Displays information about the number of LSA counts per router.

Command Default

IPv4 and unicast address prefixes

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

The various forms of the **show ospf database** command deliver information about different OSPF link-state advertisements. This command can be used to examine the link-state database (LSD) and its contents. Each router participating in an area having identical database entries pertaining to that area (with the exception of LSAs that are being flooded). Numerous options (such as **network** and **router**) are used to display portions of the database.

Task ID

Task ID	Operations
ospf	read

Examples

The following is sample output from the **show ospf database** command when no arguments or keywords are used:

```
RP/0/RP0/CPU0:router# show ospf database

OSPF Router with ID (172.20.1.11) (Process ID 1)

                Router Link States (Area 0)

 Link ID         ADV Router      Age         Seq#       Checksum Link count
 172.20.1.8      172.20.1.8      1381        0x8000010D    0xEF60   2
 172.20.1.11     172.20.1.11     1460        0x800002FE    0xEB3D   4
 172.20.1.12     172.20.1.12     2027        0x80000090    0x875D   3
 172.20.1.27     172.20.1.27     1323        0x800001D6    0x12CC   3
```

```
                Net Link States (Area 0)

Link ID         ADV Router      Age         Seq#       Checksum
172.22.1.27      172.20.1.27      1323      0x8000005B    0xA8EE
172.22.1.11      172.20.1.11      1461      0x8000005B    0x7AC

                Type-10 Opaque Link Area Link States (Area 0)

Link ID         ADV Router      Age         Seq#       Checksum Opaque ID
10.0.0.0         172.20.1.11      1461      0x800002C8    0x8483     0
10.0.0.0         172.20.1.12      2027      0x80000080    0xF858     0
10.0.0.0         172.20.1.27      1323      0x800001BC    0x919B     0
10.0.0.1         172.20.1.11      1461      0x8000005E    0x5B43     1
```

This table describes the significant fields shown in the display.

Table 43: show ospf database Field Descriptions

Field	Description
Link ID	Router ID number.
ADV Router	ID of the advertising router.
Age	Link-state age.
Seq#	Link-state sequence number (detects old or duplicate LSAs).
Checksum	Fletcher checksum of the complete contents of the LSA.
Link count	Number of interfaces detected for the router.
Opaque ID	Opaque LSA ID number.

The following is sample output from the **show ospf database** command with the **asbr-summary** keyword:

```
RP/0/RSP0RP0/CPU0:router# show ospf database asbr-summary

OSPF Router with ID (192.168.0.1) (Process ID 300)

Summary ASB Link States (Area 0.0.0.0)

  LS age: 1463
  Options: (No TOS-capability)
  LS Type: Summary Links (AS Boundary Router)
  Link State ID: 172.17.245.1 (AS Boundary Router address)
  Advertising Router: 172.17.241.5
  LS Seq Number: 80000072
  Checksum: 0x3548
  Length: 28
  Network Mask: /0
 TOS: 0  Metric: 1
```

This table describes the significant fields shown in the display.

Table 44: show ospf database asbr-summary Field Descriptions

Field	Description
OSPF Router with ID	Router ID number.
Process ID	OSPF process name.
LS age	Link-state age.
Options	Type of service options (Type 0 only).
LS Type	Link-state type.
Link State ID	Link-state ID (ASBR).
Advertising Router	ID of the advertising router.
LS Seq Number	Link-state sequence (detects old or duplicate LSAs).
Checksum	Link-state checksum (Fletcher checksum of the complete contents of the LSA).
Length	Length (in bytes) of the LSAs.
Network Mask	Network mask implemented.
TOS	Type of service.
Metric	Link-state metric.

The following is sample output from the **show ospf database** command with the **external** keyword:

```
RP/0/RP0/CPU0:router# show ospf database external

OSPF Router with ID (192.168.0.1) (Process ID 300)

                  Type-5 AS External Link States

  LS age: 280
  Options: (No TOS-capability)
  LS Type: AS External Link
  Link State ID: 172.17.0.0 (External Network Number)
  Advertising Router: 172.17.70.6
  LS Seq Number: 80000AFD
  Checksum: 0xC3A
  Length: 36
  Network Mask: 255.255.0.0
        Metric Type: 2 (Larger than any link state path)
        TOS: 0
        Metric: 1
        Forward Address: 0.0.0.0
        External Route Tag: 0
```

This table describes the significant fields shown in the display.

Table 45: show ospf database external Field Descriptions

Field	Description
OSPF Router with Router ID	Router ID number.
Process ID	OSPF process name.
LS age	Link-state age.
Options	Type of service options (Type 0 only).
LS Type	Link-state type.
Link State ID	Link-state ID (external network number).
Advertising Router	ID of the advertising router.
LS Seq Number	Link-state sequence number (detects old or duplicate LSAs).
Checksum	Link-state checksum (Fletcher checksum of the complete contents of the LSA).
Length	Length (in bytes) of the LSA.
Network Mask	Network mask implemented.
Metric Type	External type.
TOS	Type of service.
Metric	Link-state metric.
Forward Address	Forwarding address. Data traffic for the advertised destination is forwarded to this address. If the forwarding address is set to 0.0.0.0, data traffic is forwarded instead to the originator of the advertisement.
External Route Tag	External route tag, a 32-bit field attached to each external route. This tag is not used by the OSPF protocol itself.

The following is sample output from the **show ospf database** command with the **network** keyword:

```
RP/0/RP0/CPU0:router# show ospf database network

 OSPF Router with ID (192.168.0.1) (Process ID 300)

Net Link States (Area 0.0.0.0)

  LS age: 1367
  Options: (No TOS-capability)
  LS Type: Network Links
  Link State ID: 172.23.1.3 (address of Designated Router)
  Advertising Router: 192.168.0.1
  LS Seq Number: 800000E7
  Checksum: 0x1229
```

```
Length: 52
Network Mask: /24
      Attached Router: 192.168.0.1
      Attached Router: 172.23.241.5
      Attached Router: 172.23.1.1
      Attached Router: 172.23.54.5
      Attached Router: 172.23.1.5
```

This table describes the significant fields shown in the display.

Table 46: show ospf database network Field Descriptions

Field	Description
OSPF Router with ID	Router ID number.
Process ID	OSPF process name.
LS age	Link-state age.
Options	Type of service options (Type 0 only).
LS Type	Link-state type.
Link State ID	Link-state ID of the designated router.
Advertising Router	ID of the advertising router.
LS Seq Number	Link-state sequence number (detects old or duplicate LSAs).
Checksum	Link-state checksum (Fletcher checksum of the complete contents of the LSA).
Length	Length (in bytes) of the LSA.
Network Mask	Network mask implemented.
Attached Router	List of routers attached to the network, by IP address.

The following is sample output, carrying Multiprotocol Label Switching traffic engineering (MPLS TE) specification information, from the **show ospf database** command with the **opaque-area** keyword and a *link-state-id* of adv-router:

```
RP/0/RP0/CPU0:router# show ospf database opaque-area adv-router 172.20.1.12

  OSPF Router with ID (172.20.1.11) (Process ID 1)

                Type-10 Opaque Link Area Link States (Area 0)

   LS age: 224
   Options: (No TOS-capability, DC)
   LS Type: Opaque Area Link
   Link State ID: 1.0.0.0
   Opaque Type: 1
   Opaque ID: 0
   Advertising Router: 172.20.1.12
   LS Seq Number: 80000081
   Checksum: 0xF659
```

```
    Length: 132
    Fragment number : 0

      MPLS TE router ID : 172.20.1.12

      Link connected to Point-to-Point network
        Link ID : 172.20.1.11
        Interface Address : 172.21.1.12
        Neighbor Address : 172.21.1.11
        Admin Metric : 10
        Maximum bandwidth : 193000
        Maximum reservable bandwidth : 125000
        Number of Priority : 8
        Priority 0 : 125000      Priority 1 : 125000
        Priority 2 : 125000      Priority 3 : 125000
        Priority 4 : 125000      Priority 5 : 125000
        Priority 6 : 125000      Priority 7 : 100000
        Affinity Bit : 0x0

      Number of Links : 1
```

The following is the sample output from the **show ospf database opaque-area** command displaying the extended link LSA information.

```
RP/0/RP0/CPU0:router# show ospf database opaque-area 4.0.0.0
LS age: 361
  Options: (No TOS-capability, DC)
  LS Type: Opaque Area Link
  Link State ID: 8.0.0.40
  Opaque Type: 8
  Opaque ID: 40
  Advertising Router: 100.0.0.3
  LS Seq Number: 8000012e
Checksum: 0xeab4
  Length: 92

    Extended Link TLV: Length: 68
      Link-type : 2
      Link ID   : 100.0.9.4
      Link Data : 100.0.9.3

     LAN Adj sub-TLV: Length: 16
       Flags      : 0x0
       MTID       : 0
       Weight     : 0
       Neighbor ID: 100.0.0.1

       SID/Label sub-TLV: Length: 3
          SID       : 24001

     LAN Adj sub-TLV: Length: 16
       Flags      : 0x0
       MTID       : 0
       Weight     : 0
       Neighbor ID: 100.0.0.2

       SID/Label sub-TLV: Length: 3
          SID       : 24000

     Adj sub-TLV: Length: 12
        Flags     : 0x0
```

```
      MTID      : 0
      Weight    : 0

     SID/Label sub-TLV: Length: 3
       SID       : 24002
```

The following is sample output from the **show ospf database** command that displays a Type 10, Router Information LSA:

```
RP/0/RP0/CPU0:router# show ospf database opaque-area 4.0.0.0

            OSPF Router with ID (3.3.3.3) (Process ID orange)

                Type-10 Opaque Link Area Link States (Area 0)

  LS age: 105
  Options: (No TOS-capability, DC)
  LS Type: Opaque Area Link
  Link State ID: 4.0.0.0
  Opaque Type: 4
  Opaque ID: 0
  Advertising Router: 3.3.3.3
  LS Seq Number: 80000052
  Checksum: 0x34e2
  Length: 52
  Fragment number: 0

    Router Information TLV: Length: 4
    Capabilities:
      Graceful Restart Helper Capable
      Traffic Engineering enabled area
      All capability bits: 0x50000000

    PCE Discovery TLV: Length: 20
      IPv4 Address: 3.3.3.3
      PCE Scope: 0x20000000
      Compute Capabilities:
      Inter-area default (Rd-bit)
      Compute Preferences:
      Intra-area: 0  Inter-area:  0
      Inter-AS:   0  Inter-layer: 0
```

This table describes the significant fields shown in the display.

Table 47: show ospf database opaque-area Field Descriptions

Field	Description
OSPF Router with ID	Router ID number.
Process ID	OSPF process name.
LS age	Link-state age.
Options	Type of service options (Type 0 only).
LS Type	Link-state type.
Link State ID	Link-state ID.

Field	Description
Opaque Type	Opaque link-state type.
Opaque ID	Opaque ID number.
Advertising Router	ID of the advertising router.
LS Seq Number	Link-state sequence (detects old or duplicate LSAs).
Checksum	Link-state checksum (Fletcher checksum of the complete contents of the LSA).
Length	Length (in bytes) of the LSA.
Fragment number	Arbitrary value used to maintain multiple traffic engineering LSAs.
Link ID	Link ID number.
Interface Address	ID address of the interface.
Neighbor Address	IP address of the neighbor.
Admin Metric	Administrative metric value used by MPLS TE.
Maximum bandwidth	Specifies maximum bandwidth (in kbps).
Maximum reservable bandwidth	Specifies maximum reservable bandwidth (in kbps).
Number of Priority	Priority number.
Affinity Bit	Used by MPLS TE.
Router Information TLV	Router capabilities are advertised in this TLV.
Capabilities	Some router capabilities include stub router, traffic engineering, graceful restart, and graceful restart helper.
PCE Discovery TLV	PCE address and capability information is advertised in this TLV.
IPv4 Address	Configured PCE IPv4 address.
PCE Scope	Computation capabilities of the PCE.
Compute Capabilities	Compute capabilities and preferences of the PCE.
Inter-area default (RD-bit)	PCE compute capabilities such as intra-area, inter-area, inter-area default, inter-AS, inter-AS default and inter-layer.
Compute Preferences	Order or preference of path computation that includes intra-area, inter-area, inter-AS, and inter-layer preferences.

The following is sample output from the **show ospf database** command with the **router** keyword:

```
RP/0/RP0/CPU0:router# show ospf database router
```

```
OSPF Router with ID (192.168.0.1) (Process ID 300)

Router Link States (Area 0.0.0.0)

  LS age: 1176
  Options: (No TOS-capability)
  LS Type: Router Links
  Link State ID: 172.23.21.6
  Advertising Router: 172.23.21.6
  LS Seq Number: 80002CF6
  Checksum: 0x73B7
  Length: 120
  AS Boundary Router
  Number of Links: 8

 Link connected to: another Router (point-to-point)
(Link ID) Neighboring Router ID: 172.23.21.5
(Link Data) Router Interface address: 172.23.21.6
 Number of TOS metrics: 0
  TOS 0 Metrics: 2
```

This table describes the significant fields shown in the display.

Table 48: show ospf database router Field Descriptions

Field	Description
OSPF Router with ID	Router ID number.
Process ID	OSPF process name.
LS age	Link-state age.
Options	Type of service options (Type 0 only).
LS Type	Link-state type.
Link State ID	Link-state ID.
Advertising Router	ID of the advertising router.
LS Seq Number	Link-state sequence (detects old or duplicate LSAs).
Checksum	Link-state checksum (Fletcher checksum of the complete contents of the LSA).
Length	Length (in bytes) of the LSA.
AS Boundary Router	Definition of router type.
Number of Links	Number of active links.
Link ID	Link type.
Link Data	Router interface address.
TOS	Type of service metric (Type 0 only).

The following is sample output from **show ospf database** command with the **summary** keyword:

```
RP/0/RP0/CPU0:router# show ospf database summary

       OSPF Router with ID (192.168.0.1) (Process ID 300)

 Summary Net Link States (Area 0.0.0.0)

  LS age: 1401
  Options: (No TOS-capability)
  LS Type: Summary Links (Network)
  Link State ID: 172.23.240.0 (Summary Network Number)
  Advertising Router: 172.23.241.5
  LS Seq Number: 80000072
  Checksum: 0x84FF
  Length: 28
  Network Mask: /24
     TOS: 0  Metric: 1
```

This table describes the significant fields shown in the display.

Table 49: show ospf database summary Field Descriptions

Field	Description
OSPF Router with ID	Router ID number.
Process ID	OSPF process name.
LS age	Link-state age.
Options	Type of service options (Type 0 only).
LS Type	Link-state type.
Link State ID	Link-state ID (summary network number).
Advertising Router	ID of the advertising router.
LS Seq Number	Link-state sequence (detects old or duplicate LSAs).
Checksum	Link-state checksum (Fletcher checksum of the complete contents of the LSA).
Length	Length (in bytes) of the LSA.
Network Mask	Network mask implemented.
TOS	Type of service.
Metric	Link-state metric.

The following is sample output from **show ospf database** command with the **database-summary** keyword:

```
RP/0/RP0/CPU0:router# show ospf database database-summary
```

```
            OSPF Router with ID (172.19.65.21) (Process ID 1)

  Area 0 database summary
    LSA Type      Count    Delete   Maxage
    Router        2        0        0
    Network       1        0        0
    Summary Net   2        0        0
    Summary ASBR  0        0        0
    Type-7 Ext    0        0        0
    Opaque Link   0        0        0
    Opaque Area   0        0        0
    Subtotal      5        0        0

  Process 1 database summary
    LSA Type      Count    Delete   Maxage
    Router        2        0        0
    Network       1        0        0
    Summary Net   2        0        0
    Summary ASBR  0        0        0
    Type-7 Ext    0        0        0
    Opaque Link   0        0        0
    Opaque Area   0        0        0
    Type-5 Ext    2        0        0
    Opaque AS     0        0        0
    Total         7        0        0
```

The **show ospf database database-summary** with the **detail** keyword displays the number of LSA counts per router:

```
Router#show ospf database database-summary detail
Tue Dec  6 19:20:34.090 IST

            OSPF Router with ID (192.168.0.1) (Process ID 1)

Router 192.168.0.4 LSA summary
  LSA Type      Count    Delete   Maxage
  Router        0        0        0
  Network       0        0        0
  Summary Net   0        0        0
  Summary ASBR  0        0        0
  Type-5 Ext    697      0        0
  Type-7 Ext    0        0        0
  Opaque Link   0        0        0
  Opaque Area   0        0        0
  Opaque AS     6        0        0
  Total         703      0        0

Router 192.168.0.1 LSA summary
  LSA Type      Count    Delete   Maxage
  Router        1        0        0
  Network       0        0        0
  Summary Net   0        0        0
  Summary ASBR  0        0        0
  Type-5 Ext    0        0        0
  Type-7 Ext    0        0        0
  Opaque Link   0        0        0
  Opaque Area   64       0        0
  Opaque AS     0        0        0
  Total         65       0        0

Router 192.168.0.2 LSA summary
```

```
  LSA Type      Count    Delete   Maxage
  Router        1        0        0
  Network       0        0        0
  Summary Net   21       0        0
  Summary ASBR  2        0        0
  Type-5 Ext    0        0        0
  Type-7 Ext    0        0        0
  Opaque Link   0        0        0
  Opaque Area   21       0        0
  Opaque AS     0        0        0
  Total         45       0        0

Router 192.168.0.6 LSA summary
  LSA Type      Count    Delete   Maxage
  Router        1        0        0
  Network       0        0        0
  Summary Net   21       0        0
  Summary ASBR  2        0        0
  Type-5 Ext    0        0        0
  Type-7 Ext    0        0        0
  Opaque Link   0        0        0
  Opaque Area   19       0        0
  Opaque AS     0        0        0
  Total         43       0        0

Router 192.168.0.3 LSA summary
  LSA Type      Count    Delete   Maxage
  Router        0        0        0
  Network       0        0        0
  Summary Net   0        0        0
  Summary ASBR  0        0        0
  Type-5 Ext    7        0        0
  Type-7 Ext    0        0        0
  Opaque Link   0        0        0
  Opaque Area   0        0        0
  Opaque AS     6        0        0
  Total         13       0        0
```

The **show ospf database database-summary** with the **adv-router***router ID* keyword displays the router information and the LSAs received from a particular router:

```
Router#show ospf database database-summary adv-router 192.168.0.4
Tue Dec  6 19:21:04.549 IST


            OSPF Router with ID (192.168.0.1) (Process ID 1)

Router 192.168.0.4 LSA summary
  LSA Type      Count    Delete   Maxage
  Router        0        0        0
  Network       0        0        0
  Summary Net   0        0        0
  Summary ASBR  0        0        0
  Type-5 Ext    697      0        0
  Type-7 Ext    0        0        0
  Opaque Link   0        0        0
  Opaque Area   0        0        0
  Opaque AS     6        0        0
  Total         703      0        0
```

This table describes the significant fields shown in the display.

Table 50: show ospf database database-summary Field Descriptions

Field	Description
LSA Type	Link-state type.
Count	Number of advertisements in that area for each link-state type.
Delete	Number of LSAs that are marked “Deleted” in that area.
Maxage	Number of LSAs that are marked “Maxaged” in that area.

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*] [**vrf** {*vrf-name* | **all**}] [*area-id*] **flood-list** [*type interface-path-id*]

Syntax Description

process-name	(Optional) OSPF process name that uniquely identifies an OSPF routing process. The process name is any alphanumeric string no longer than 40 characters. If this argument is included, only information for the specified routing process is included.
vrf	(Optional) Specifies an OSPF VPN routing and forwarding (VRF) instance.
vrf-name	(Optional) Name of the OSPF VRF. The *vrf-name* argument can be specified as an arbitrary string. The strings "default" and "all" are reserved VRF names.
all	(Optional) Specifies all OSPF VRF instances.
area-id	(Optional) Area number used to define the particular area.
type	Interface type. For more information, use the question mark (?) online help function.
interface-path-id	Physical interface or virtual interface. **Note** Use the **show interfaces** command to see a list of all interfaces currently configured on the router. For more information about the syntax for the router, use the question mark (?) online help function.

Command Default

All interfaces

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

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 HundredGigE 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 51: 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 strict-mode information use the **show ospf interface** command in XR EXEC mode.

show ospf interface [**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

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

Task ID

Task ID	Operation
bgp	read
ospf	read

show ospf interface

The following example shows the output from the **show ospf interface** command.

```
RP/0/0/CPU0:rt2(config-ospf-ar-if)# show ospf interface HundredGigE 0/2/0/0
Sun Feb 15 12:17:35.072 IST

HundredGigE 0/2/0/0 is up, line protocol is up
  Internet Address 10.1.1.2/24, Area 0
  Process ID 1, Router ID 2.2.2.2, Network Type BROADCAST, Cost: 1
  Transmit Delay is 1 sec, State DR, Priority 1, MTU 1500, MaxPktSz 1500
  BFD enabled, BFD interval 150 msec, BFD multiplier 3, Mode: Strict
  Designated Router (ID) 2.2.2.2, Interface address 10.1.1.2
  No backup designated router on this network
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
    Hello due in 00:00:07:358
  Index 1/1, 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 1
Neighbor Count is 1, Adjacent neighbor count is 0
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*] [**vrf** {*vrf-name* | **all**}] [*area-id*] [*type interface-path-id*] **mpls traffic-eng** {**link** | **fragment**}

Syntax Description

process-name	(Optional) OSPF process name that uniquely identifies an OSPF routing process. The process name is any alphanumeric string no longer than 40 characters. If this argument is included, only information for the specified routing process is included.
vrf *vrf-name* **all**	(Optional) Specifies an OSPF VPN routing and forwarding (VRF) instance. The *vrf-name* argument can be specified as an arbitrary string. The strings "default" and "all" are reserved VRF names.
area-id	(Optional) Area number used to define the particular area.
type	Interface type. For more information, use the question mark (?) online help function.
interface-path-id	Physical interface or virtual interface. **Note** Use the **show interfaces** command to see a list of all interfaces currently configured on the router. For more information about the syntax for the router, use the question mark (?) online help function.
link	Provides detailed information about the links over which traffic engineering is supported on the local router.
fragment	Provides detailed information about the traffic engineering fragments on the local router.

Command Default

All links or fragments

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

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 52: show ospf mpls traffic-eng link Field Descriptions

Field	Description
Link ID	Link type.
Interface address	IP address of the interface.
Neighbor address	IP address of the neighbor.
Admin Metric	Administrative distance metric value used by Multiprotocol Label Switching traffic engineering (MPLS TE).
Maximum bandwidth	Bandwidth capacity of the link (in kbps).
Maximum global pool reservable bandwidth	Maximum amount of bandwidth that is available for reservation in the global pool.
Maximum sub pool reservable bandwidth	Maximum amount of bandwidth that is available for reservation in the subpool.
Number of Priority	Priority number.
Global pool unreserved BW	Amount of unreserved bandwidth that is available in the global pool.
Sub pool unreserved BW	Amount of unreserved bandwidth that is available in the subpool.
Affinity Bit	Used by MPLS TE. Attribute values required for links carrying this tunnel. A 32-bit dotted-decimal number. Valid values are from 0x0 to 0xFFFFFFFF, representing 32 attributes (bits), where the value of an attribute is 0 or 1.

The following is sample output from the **show ospf mpls traffic-eng** command when the **fragment** keyword is specified:

```
RP/0/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 53: show ospf mpls traffic-eng fragment Field Descriptions

Field	Description
Area instance	Number of times traffic engineering information or any link changed.
Link instance	Number of times any link changed.
Link ID	Link type.
Interface address	IP address of the interface.
Neighbor address	IP address of the neighbor.
Admin Metric	Administrative distance metric value used by MPLS TE.
Maximum bandwidth	Bandwidth capacity of the link (in kbps).

Field	Description
Maximum global pool reservable bandwidth	Maximum amount of bandwidth that is available for reservation in the global pool.
Maximum sub pool reservable bandwidth	Maximum amount of bandwidth that is available for reservation in the subpool.
Number of Priority	Priority number.
Global pool unreserved BW	Amount of unreserved bandwidth that is available in the global pool.
Sub pool unreserved BW	Amount of unreserved bandwidth that is available in the subpool.
Affinity Bit	Used by MPLS TE. Attribute values required for links carrying this tunnel. A 32-bit dotted-decimal number. Valid values are from 0x0 to 0xFFFFFFFF, representing 32 attributes (bits), where the value of an attribute is 0 or 1.

show ospf message-queue

To display the information about the queue dispatch values, peak lengths, and limits, use the **show ospf message-queue** command in 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 mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

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 54: show ospf message-queue Field Descriptions

Field	Description
Hello Input Queue	This section provides statistics on the number of events and incoming packets processed in the Hello (incoming packet) thread of the OSPF process.
Router Message Queue	This section provides statistics on the events and messages processed in the Router (primary) thread of the OSPF process.
TE Message Queue	This section provides statistics on traffic-engineering events and messages received by OSPF from TE (the te_control process). These events are processed in the Router thread of the OSPF process.
Number of Dlink errors	The number of enqueuing or dequeuing errors seen across all the linked-lists in the OSPF process.

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*] [**vrf** {*vrf-name* | **all**}] [*area-id*] **neighbor** [[*type interface-path-id*] [*neighbor-id*] [**detail**] | **area-sorted**]

Syntax Description

process-name	(Optional) Name that uniquely identifies an OSPF routing process. The process name is defined by the **router ospf** command. If this argument is included, only information for the specified routing process is displayed.
vrf *vrf-name* **all**	(Optional) Specifies an OSPF VPN routing and forwarding (VRF) instance. The *vrf-name* argument can be specified as an arbitrary string. The strings "default" and "all" are reserved VRF names.
area-id	(Optional) Area ID. If you do not specify an area, all areas are displayed.
type	Interface type. For more information, use the question mark (?) online help function.
interface-path-id	Physical interface or virtual interface. **Note** Use the **show interfaces** command to see a list of all interfaces currently configured on the router. For more information about the syntax for the router, use the question mark (?) online help function.
neighbor-id	(Optional) Neighbor ID.
detail	(Optional) Displays all neighbors given in detail (lists all neighbors).
area-sorted	(Optional) Specifies that all neighbors are grouped by area.

Command Default

All neighbors

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

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      HundredGigE 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       HundredGigE 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     HundredGigE 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     HundredGigE 0/3/0/3
    Neighbor is up for 18:45:27
```

This table describes the significant fields shown in the display.

Table 55: 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 HundredGigE 0/3/0/2
    Neighbor priority is 1, State is FULL, 6 state changes
    DR is 0.0.0.0 BDR is 0.0.0.0
    Options is 0x2
    Dead timer due in 0:00:32
    Neighbor is up for 18:45:30
    Number of DBD retrans during last exhange 0
    Index 1/1, retransmission queue length 0, number of retransmission 0
    First 0x0(0)/0x0(0) Next 0x0(0)/0x0(0)
    Last retransmission scan length is 0, maximum is 0
    Last retransmission scan time is 0 msec, maximum 0 msec
 Neighbor 192.168.199.137, interface address 192.168.48.189
```

```
    In the area 0.0.0.0 via interface HundredGigE 0/3/0/3
    Neighbor priority is 5, State is FULL, 6 state changes
    Options is 0x2
    Dead timer due in 0:00:32
    Neighbor is up for 18:45:30
    Number of DBD retrans during last exhange 0
    Index 1/1, retransmission queue length 0, number of retransmission 0
    First 0x0(0)/0x0(0) Next 0x0(0)/0x0(0)
    Last retransmission scan length is 0, maximum is 0
    Last retransmission scan time is 0 msec, maximum 0 msec

Total neighbor count: 2
```

This table describes the significant fields shown in the display.

Table 56: 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 HundredGigE 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 HundredGigE 0/3/0/2
    Neighbor priority is 1, State is FULL, 6 state changes
    DR is 0.0.0.0 BDR is 0.0.0.0
    Options is 0x2
```

```
    Dead timer due in 0:00:32
    Neighbor is up for 18:45:30
    Number of DBD retrans during last exhange 0
    Index 1/1, retransmission queue length 0, number of retransmission 0
    First 0x0(0)/0x0(0) Next 0x0(0)/0x0(0)
    Last retransmission scan length is 0, maximum is 0
    Last retransmission scan time is 0 msec, maximum 0 msec

Total neighbor count: 1
```

This table describes the significant fields shown in the display.

Table 57: show ospf neighbor HundredGigE 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 HundredGigE 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       HundredGigE 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     HundredGigE 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     HundredGigE 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 58: show ospf neighbor HundredGigE 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 HundredGigE 0/3/0/2
    Neighbor priority is 1, State is FULL, 6 state changes
    DR is 0.0.0.0 BDR is 0.0.0.0
    Options is 0x2
    Dead timer due in 0:00:32
    Neighbor is up for 18:45:30
    Number of DBD retrans during last exhange 0
    Index 1/1, retransmission queue length 0, number of retransmission 0
    First 0x0(0)/0x0(0) Next 0x0(0)/0x0(0)
    Last retransmission scan length is 0, maximum is 0
    Last retransmission scan time is 0 msec, maximum 0 msec

Total neighbor count: 1
```

```
Neighbor 10.1.1.1, interface address 192.168.13.1
   In the area 0 via interface HundredGigE 0/3/0/1
   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 HundredGigE 0/3/0/2
   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 59: 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.

Field	Description
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.

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*] [**vrf** {*vrf-name* | **all**}] [*area-id*] **request-list** [*type interface-path-id*] [*neighbor-id*]

Syntax Description

process-name	(Optional) Name that uniquely identifies an OSPF routing process. The process name is defined by the **router ospf** command. If this argument is included, only information for the specified routing process is displayed.
vrf	(Optional) Specifies an OSPF VPN routing and forwarding (VRF) instance.
vrf-name	(Optional) Name of the OSPF VRF. The *vrf-name* argument can be specified as an arbitrary string. The strings “default” and “all” are reserved VRF names.
all	(Optional) Specifies all OSPF VRF instances.
area-id	(Optional) Area ID. If you do not specify an area, all areas are displayed.
type	Interface type. For more information, use the question mark (?) online help function.
i *nterface-path-id*	Physical interface or virtual interface. Use the **show interfaces** command to see a list of all interfaces currently configured on the router. For more information about the syntax for the router, use the question mark (?) online help function.
neighbor-id	(Optional) IP address of the OSPF neighbor.

Command Default

All neighbors

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

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 HundredGigE 3/0/0/0

Request Lists for OSPF pagent

 Neighbor 10.0.124.4, interface HundredGigE 3/0/0/0 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 60: show ospf request-list 10.0.124.4 HundredGigE 3/0/0/0 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.

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*] [**vrf** {*vrf-name* | **all**}] [*area-id*] **retransmission-list** [*type interface-path-id*] [*neighbor-id*]

Syntax Description

process-name	(Optional) Name that uniquely identifies an OSPF routing process. The process name is defined by the **router ospf** command. If this argument is included, only information for the specified routing process is displayed.
vrf *vrf-name* **all**	(Optional) Specifies an OSPF VPN routing and forwarding (VRF) instance. The *vrf-name* argument can be specified as an arbitrary string. The strings "default" and "all" are reserved VRF names.
area-id	(Optional) Area ID. If you do not specify an area, all areas are displayed.
type	Interface type. For more information, use the question mark (?) online help function.
interface-path-id	Physical interface or virtual interface. **Note** Use the **show interfaces** command to see a list of all interfaces currently configured on the router. For more information about the syntax for the router, use the question mark (?) online help function.
neighbor-id	(Optional) IP address of the OSPF neighbor.

Command Default

All neighbors

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

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 HundredGigE 3/0/0/0

 Neighbor 10.0.124.4, interface HundredGigE 3/0/0/0 address 10.3.1.2
```

This table describes the significant fields shown in the display.

Table 61: 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.

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*] [**vrf** {*vrf-name* | **all**}] **routes** [**connected** | **external** | **local**] [*prefix mask*] [*prefix/length*] **[multicast-intact] [backup-path]**

Syntax Description

process-name	(Optional) Name that uniquely identifies an OSPF routing process. The process name is defined by the **router ospf** command. If this argument is included, only information for the specified routing process is displayed.
vrf *vrf-name* **all**	(Optional) Specifies an OSPF VPN routing and forwarding (VRF) instance. The *vrf-name* argument can be specified as an arbitrary string. The strings "default" and "all" are reserved VRF names.
connected	(Optional) Displays connected routes.
external	(Optional) Displays routes redistributed from other protocols.
local	(Optional) Displays the local routes redistributed from the Routing Information Base (RIB).
prefix	(Optional) IP prefix, which limits output to a specific route. If the *prefix* argument is specified, either the *length* or *mask* argument is required.
mask	(Optional) IP address mask.
/ *length*	(Optional) Prefix length, which can be indicated as a slash (/) and number. For example, /8 indicates that the first eight bits in the IP prefix are network bits. If *length* is used, the slash is required.
multicast-intact	(Optional) Displays multicast intact paths.
backup-path	(Optional) Displays fast-reroute backup path information.

Command Default

All route types

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 7.0.12	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.

When you redistribute BGP routes into OSPF, the OSPF external topology table may display the output interface as NULLIFHNDL. This occurs because BGP routes are recursive, relying on an IP next-hop instead of a direct output interface. You can ignore NULLIFHNDL in this context; it does not indicate a functional error or synchronization issue.

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, O IA - Inter area
     O E1 - External type 1, O E2 - External type 2
     O N1 - NSSA external type 1, O N2 - NSSA external type 2

O    2.2.2.2/32, metric 3
     10.1.0.2, from 2.2.2.2, via HundredGigE 0/0/0/7, 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 HundredGigE 0/0/0/3, protected bitmap 0x1
            Attribues: Metric: 104, SRLG Disjoint
```

This table describes the significant fields shown in the display.

Table 62: show ospf route Field Descriptions

Field	Description
O	OSPF route.
E	External Type 1 or 2 route.
N	NSSA Type 1 or 2
2.2.2.2/32	Network and subnet mask to which the local router has a route.
metric	Cost to reach network 10.3.1.0.
10.1.0.2	Next-hop router on the path to network 10.3.1.0.
from 2.2.2.2	Router ID 172.16.10.1 is the router that advertised this route.
via [Interface]	Packets destined for the given prefix (10.3.1.0/24) are sent over the specified interface (GigabitEthernet interface 0/0/0/7). **Note** If you redistribute external routes from BGP, you might see NULLIFHNDL in this field.

Field	Description
Backup path	Indicates the topology independent loop-free alternate backup path. Here, the backup path uses the P node 4.4.4.4.

Examples

The following is sample output from the **show ospf routes** command:

```
RP/0/RP0/CPU0:router# show ospf routes

Topology Table for ospf 1 with ID 10.3.4.2

Codes:O - Intra area, O IA - Inter area
       O E1 - External type 1, O E2 - External type 2
       O N1 - NSSA external type 1, O N2 - NSSA external type 2

O E2 10.3.1.0/24, metric 1
       10.3.4.1, from 172.16.10.1, via HundredGigE 0/1/0/1
O    10.3.4.0/24, metric 1562
       10.3.4.2, directly connected, via HundredGigE 0/1/0/1
O E2 10.1.0.0/16, metric 1
       10.3.4.1, from 172.16.10.1, via HundredGigE 0/1/0/1
O IA 10.10.10.0/24, metric 1572
       10.3.4.1, from 172.16.10.1, via HundredGigE 0/1/0/1
O E2 130.10.10.0/24, metric 20
       10.3.4.1, from 172.16.10.1, via HundredGigE 0/1/0/1
```

This table describes the significant fields shown in the display.

Table 63: show ospf route Field Descriptions

Field	Description
O	OSPF route.
E	External Type 1 or 2 route.
N	NSSA Type 1 or 2
10.3.1.0/24	Network and subnet mask to which the local router has a route.
metric	Cost to reach network 10.3.1.0.
10.3.4.1	Next-hop router on the path to network 10.3.1.0.
from 172.16.10.1	Router ID 172.16.10.1 is the router that advertised this route.
via GigabitEthernet 0/1/0/1	Packets destined for the given prefix (10.3.1.0/24) are sent over GigabitEthernet interface 0/1/0/1.

This table describes the significant fields shown in the display.

Table 64: show ospf route Field Descriptions

Field	Description
O	OSPF route.
E2	External Type 2 route.
10.3.1.0/24	Network and subnet mask to which the local router has a route.
metric 1	Cost to reach network 10.3.1.0.
10.3.4.1	Next-hop router on the path to network 10.3.1.0.
from 172.16.10.1	Router ID 172.16.10.1 is the router that advertised this route.
via POS 0/1/0/1	Packets destined for the given prefix (10.3.1.0/24) are sent over POS interface 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, O IA - Inter area
       O E1 - External type 1, O E2 - External type 2
       O N1 - NSSA external type 1, O N2 - NSSA external type 2

O    10.1.5.0/24, metric 1562
       10.1.5.14, directly connected, via HundredGigE 0/3/0/3
O IA 21.0.0.0/24, metric 1572
       10.1.5.12, from 172.23.54.12, via HundredGigE 0/3/0/3
O    10.0.0.0/24, metric 10
       10.0.0.12, directly connected, via HundredGigE 0/2/0/3
```

This table describes the significant fields shown in the display.

Table 65: 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.

Field	Description
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 66: 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.
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.

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*] [**vrf** {*vrf-name* | **all**}] **sham-links**

Syntax Description

process-name	(Optional) Name that uniquely identifies an OSPF routing process. The process name is defined by the **router ospf** command. If this argument is included, only information for the specified routing process is displayed.
vrf	(Optional) Specifies an OSPF VPN routing and forwarding (VRF) instance.
vrf-name	(Optional) Name of the OSPF VRF. The *vrf-name* argument can be specified as an arbitrary string. The strings “default” and “all” are reserved VRF names.
all	(Optional) Specifies all OSPF VRF instances.

Command Default

No default behavior or values

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

Use the **show ospf sham-links** command to display OSPF sham-link information.

Task ID

Task ID	Operations
ospf	read

Examples

The following is sample output from the **show ospf sham-links** command:

```
RP/0/RP0/CPU0:router# show ospf 1 vrf vrf_1 sham-links

Sham Links for OSPF 1, VRF vrf_1

Sham Link OSPF_SL0 to address 10.0.0.3 is up
Area 0, source address 10.0.0.1
IfIndex = 185
Run as demand circuit
DoNotAge LSA allowed., Cost of using 1
Transmit Delay is 1 sec, State POINT_TO_POINT,
Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
Hello due in 00:00:04
Adjacency State FULL (Hello suppressed)
Number of DBD retrans during last exchange 0
Index 2/2, retransmission queue length 0, number of retransmission 0
```

```
First 0(0)/0(0) Next 0(0)/0(0)
Last retransmission scan length is 0, maximum is 0
Last retransmission scan time is 0 msec, maximum is 0 msec
Keychain-based authentication enabled
Key id used is 2
```

This table describes the significant fields shown in the display.

Table 67: show ospf sham-links Field Descriptions

Field	Description
Sham Link OSPF_SL0 to address	Address of the destination endpoint of the sham link.
IfIndex	ifindex associated with the sham link.
Run as demand circuit	Sham link is treated as a demand circuit.
DoNotAge LSA allowed	DoNotAge LSAs are allowed to be flooded over the sham link.
Cost of using	Sham-link cost.
Transmit Delay	Sham-link transmit delay.
State	Sham-link interface state.
Timer intervals configured	Various sham-link interface-related timers.
Hello due in	Time before the next Hello is sent over the sham link.
Adjacency State	State of the adjacency with the neighbor over the sham link.
Number of DBD retrans during last exchange	Number of DBD retransmissions during the last exchange over the sham link.
Index	Area flood index.
retransmission queue length	Retransmission queue length on the sham link.
number of retransmission	Number of retransmissions over the sham-link interface.
First	First flood information.
Next	Next flood information.
Last retransmission scan length is	Last retransmission scan length on the sham-link interface.
maximum is	Maximum retransmission scan length on the sham-link interface.
Last retransmission scan time is	Last retransmission scan time on the sham-link interface.
maximum is 0 msec	Maximum retransmission scan time on the sham-link interface.
Keychain-based authentication enabled	Keychain-based authentication is enabled.
Key id used is	Key ID used.

show ospf summary-prefix

To display Open Shortest Path First (OSPF) aggregated summary address information, use the **show ospf summary-prefix** command in XR EXEC mode.

show ospf [*process-name*] [**vrf** {*vrf-name* | **all**}] **summary-prefix**

Syntax Description

process-name	(Optional) Name that uniquely identifies an OSPF routing process. The process name is defined by the **router ospf** command. If this argument is included, only information for the specified routing process is displayed.
vrf *vrf-name* **all**	(Optional) Specifies an OSPF VPN routing and forwarding (VRF) instance. The *vrf-name* argument can be specified as an arbitrary string. The strings "default" and "all" are reserved VRF names.

Command Default

All summary prefixes

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

Use the **show ospf summary-prefix** command if you configured summarization of external routes with the **summary-prefix** command and you want to display configured summary addresses.

Task ID

Task ID	Operations
ospf	read

Examples

The following is sample output from the **show ospf summary-prefix** command:

```
RP/0/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 68: 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.

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*] [**vrf** {*vrf-name* | **all**}] **virtual-links**

Syntax Description

process-name	(Optional) Name that uniquely identifies an OSPF routing process. The process name is defined by the **router ospf** command. If this argument is included, only information for the specified routing process is displayed.
vrf *vrf-name* **all**	(Optional) Specifies an OSPF VPN routing and forwarding (VRF) instance. The *vrf-name* argument can be specified as an arbitrary string. The strings "default" and "all" are reserved VRF names.

Command Default

All virtual links

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

Use the **show ospf virtual-links** command to display useful information for debugging OSPF routing operations.

Task ID

Task ID	Operations
ospf	read

Examples

The following is sample output from the **show ospf virtual-links** command:

```
RP/0/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 69: show ospf virtual-links Field Descriptions

Field	Description
Virtual Link to router 172.31.101.2 is up	OSPF neighbor and whether the link to that neighbor is up or down.
Transit area 0.0.0.1	Transit area through which the virtual link is formed.
via interface GigabitEthernet 0/3/0/0	Interface through which the virtual link is formed.
Cost of usingusing 10	Cost of reaching the OSPF neighbor through the virtual link.
Transmit Delay is 1 sec	Transmit delay (in seconds) on the virtual link.
State POINT_TO_POINT	State of the OSPF neighbor.
Timer intervals	Various timer intervals (in seconds) configured for the link.
Hello due in 0:00:08	When the next hello message is expected from the neighbor (in hh:mm:ss).
Adjacency State FULL	Adjacency state between the neighbors.

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**] [**all***protocol*]

Syntax Description

afi-all	(Optional) Specifies all address families.
ipv4	(Optional) Specifies an IPv4 address family.
ipv6	(Optional) Specifies an IPv6 address family.
all	(Optional) Specifies all protocols for a given address family.
protocol	(Optional) Specifies a routing protocol. For the IPv4 address family, the options are: • **bgp** • **isis** • **ospf** • **rip** For the IPv6 address family, the options are: • **bgp** • **isis** • **ospfv3**

Command Default

No default behavior or value

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

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 HundredGigE  0/3/0/0
  !
  interface HundredGigE 0/3/0/1
  !
  interface HundredGigE 0/3/0/2
  !
  interface HundredGigE 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
    HundredGigE 0/3/0/3
    HundredGigE 0/3/0/2
    HundredGigE 0/3/0/1
    HundredGigE 0/3/0/0
    Loopback2
    Loopback0
```

This table describes the significant fields shown in the display.

Table 70: show protocols ospf Field Descriptions

Field	Description
Router Id	ID of the router for this configuration.
Distance	Administrative distance of OSPF routes relative to routes from other protocols.
Non-Stop Forwarding	Status of nonstop forwarding.
Redistribution	Lists the protocols that are being redistributed.
Area	Information about the current area including list of interfaces and the status of Multiprotocol Label Switching traffic engineering (MPLS TE).

snmp context (OSPF)

To specify an SNMP context for an OSPF instance, use the **snmp context** command in XR Config mode or in VRF configuration mode. To remove the SNMP context, use the **no** form of this command.

snmp context *context_name*

Syntax Description

context_name	Specifies name of the SNMP context for OSPF instance.

Command Default

SNMP context is not specified.

Command Modes

XR Config mode

VRF configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

The snmp-server commands need to be configured to perform SNMP request for the OSPF instance. Refer *SNMP Server Commands* module in *System Management Command Reference* for information on using the **snmp-server** commands.

Note To map an SNMP context with a protocol instance, topology or VRF entity, use the **snmp-server context mapping** command. However, the **feature** option of this command does not work with OSPF protocol.

Task ID

Task ID	Operation
ospf	read, write

This example shows how to configure an SNMP context *foo* for OSPF instance *100*:

```
RP/0/RP0/CPU0:router#configure
RP/0/RP0/CPU0:router(config)#router ospf 100
RP/0/RP0/CPU0:router(config-ospf)#snmp context foo
```

This example shows how to configure **snmp-server** commands to be used with the **snmp context** command:

```
RP/0/RP0/CPU0:router(config)#snmp-server host 10.0.0.2 traps version 2c public udp-port
1620
RP/0/RP0/CPU0:router(config)#snmp-server community public RW
RP/0/RP0/CPU0:router(config)#snmp-server contact foo
```

```
RP/0/RP0/CPU0:router(config)#snmp-server community-map public context foo
```

This is a sample SNMP context configuration for OSPF instance *100*:

```
snmp-server host 10.0.0.2 traps version 2c public udp-port 1620
snmp-server community public RW
snmp-server contact foo

snmp-server community-map public context foo


router ospf 100
 router-id 2.2.2.2
 bfd fast-detect
 nsf cisco
snmp context foo
 area 0
  interface Loopback1
  !
 !
 area 1
  interface  0/2/0/1
   demand-circuit enable
  !
  interface  0/3/0/0
  !
  interface  0/3/0/1
  !
 !
!
```

snmp trap (OSPF)

To enable SNMP trap for an OSPF instance, use the **snmp trap** command in VRF configuration mode. To disable SNMP trap for the OSPF instance, use the **no** form of this command.

snmp trap

Syntax Description

This command has no keywords or arguments.

Command Default

Disabled.

Command Modes

VRF configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

Task ID

Task ID	Operation
ospf	read, write

This example shows how to enable SNMP trap for OSPF instance *100* under VRF *vrf-1*:

```
RP/0/RP0/CPU0:router#configure
RP/0/RP0/CPU0:router(config)#router ospf 100
RP/0/RP0/CPU0:router(config-ospf)#vrf vrf-1
RP/0/RP0/CPU0:router(config-ospf-vrf)#snmp trap
```

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 XR Config 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*

Syntax Description

window-size	Specifies the trap rate limit sliding window size.
max-num-traps	Specifies the maximum number of traps sent in window time.

Command Default

None

Command Modes

XR Config mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

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*

Syntax Description

route-policy *policy-name*	Specifies the route policy to apply to OSPFv2 prefix prioritization.

Note
If SPF prefix prioritization is configured, /32 prefixes are no longer preferred by default. To retain the /32 prefixes in higher-priority queues, define the route-policy accordingly.

Command Default

SPF prefix prioritization is disabled.

Command Modes

OSPF router configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

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
```

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**]

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 7.0.12	This command was introduced.

Usage Guidelines

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 HundredGigE 0/3/0/3
```

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*]

Syntax Description

address	Summary address designated for a range of addresses.
mask	IP subnet mask used for the summary route.
not-advertise	(Optional) Suppresses summary routes that match the address and mask pair from being advertised.
tag *tag*	(Optional) Tag value that can be used as a "match" value for controlling redistribution through route policies.

Command Default

When this command is not used, specific addresses are created for each route from another route source being distributed into the OSPF protocol.

Command Modes

Router configuration

VRF configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

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
```

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*

Syntax Description

seconds	Interval (in seconds) at which LSAs are grouped and refreshed, checksummed, or aged. Range is 10 seconds to 1800 seconds.

Command Default

seconds : 240 seconds

Command Modes

Router configuration

VRF configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

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*

Syntax Description

milliseconds	Minimum interval (in milliseconds) between accepting same LSA. Range is 0 to 600000 milliseconds.

Command Default

milliseconds : 100 milliseconds

Command Modes

Router configuration

VRF configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

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 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*

Syntax Description

start-interval	Delay to generate first occurance of LSA in milliseconds. Range is 0 to 600000 milliseconds.
hold-interval	Minimum delay between originating the same LSA in milliseconds. Range is 1 to 600000 milliseconds.
max-interval	Maximum delay between originating the same LSA in milliseconds. Range is 1 to 600000 milliseconds.

Command Default

start-interval : 50 milliseconds

hold-interval : 200 milliseconds

max-interval : 5000 milliseconds

Command Modes

Router configuration

VRF configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

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 msecs
 Minimum hold time between two consecutive SPFs 10000 msecs
 Maximum wait time between two consecutive SPFs 10000 msecs
 Initial LSA throttle delay 500 msecs
 Minimum hold time for LSA throttle 1000 msecs
 Maximum wait time for LSA throttle 90000 msecs
 Minimum LSA interval 1000 msecs. Minimum LSA arrival 1 secs
 Maximum number of configured interfaces 255
 Number of external LSA 0. Checksum Sum 00000000
 Number of opaque AS LSA 0. Checksum Sum 00000000
 Number of DCbitless external and opaque AS LSA 0
 Number of DoNotAge external and opaque AS LSA 0
 Number of areas in this router is 2. 2 normal 0 stub 0 nssa
 External flood list length 0
 Non-Stop Forwarding enabled
    Area BACKBONE(0) (Inactive)
        Number of interfaces in this area is 2
        SPF algorithm executed 8 times
        Number of LSA 2. Checksum Sum 0x01ba83
        Number of opaque link LSA 0. Checksum Sum 00000000
        Number of DCbitless LSA 0
        Number of indication LSA 0
        Number of DoNotAge LSA 0
        Flood list length 0
    Area 1
        Number of interfaces in this area is 1
        SPF algorithm executed 9 times
        Number of LSA 2. Checksum Sum 0x0153ea
        Number of opaque link LSA 0. Checksum Sum 00000000
        Number of DCbitless LSA 0
        Number of indication LSA 0
        Number of DoNotAge LSA 0
        Flood list length 0
```

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*

Syntax Description

spf-start	Initial SPF schedule delay (in milliseconds). Range is 1 to 600000 milliseconds.
spf-hold	Minimum hold time (in milliseconds) between two consecutive SPF calculations. Range is 1 to 600000 milliseconds.
spf-max-wait	Maximum wait time (in milliseconds) between two consecutive SPF calculations. Range is 1 to 600000 milliseconds.

Command Default

spf-start:50 milliseconds

spf-hold: 200 milliseconds

spf-max-wait: 5000 milliseconds

Command Modes

Router configuration

VRF configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

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*

Syntax Description

seconds	Time (in seconds) required to send a link-state update. Range is 1 to 65535 seconds.

Command Default

seconds: 1 second

Command Modes

XR Config mode

Area configuration

Interface configuration

Virtual-link configuration

VRF configuration

Multi-area configuration

Sham-link configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

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 HundredGigE 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 HundredGigE 0/3/0/0
RP/0/RP0/CPU0:router(config-ospf-ar-if)# transmit-delay 3
```

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*

Syntax Description

router-id	Router ID associated with the virtual link neighbor. The router ID appears in the **show ospf** command display. The router ID can be any 32-bit router ID value specified in four-part, dotted-decimal notation.

Command Default

No virtual links are defined.

Command Modes

Area configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

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
```

vrf (OSPF)

To configure an Open Shortest Path First (OSPF) VPN routing and forwarding (VRF) instance, use the **vrf** command in XR Config mode. To terminate an OSPF VRF, use the **no** form of this command.

vrf *vrf-name*

Syntax Description

vrf-name	Identifier of an OSPF VRF. The *vrf-name* argument can be specified as an arbitrary string. The strings "default" and "all" are reserved VRF names.

Command Default

No OSPF VRF is defined.

Command Modes

Router configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

Use the **vrf** command to explicitly configure a VRF. Commands configured under the VRF configuration mode (such as the **interface** [OSPF] and **authentication** commands) are automatically bound to that VRF.

To modify or remove the VRF, the *vrf-id* argument format must be the same as the format used when creating the area.

Note To remove the specified VRF from the router configuration, use the **no vrf** *vrf-id* command. The **no vrf** *vrf-id* command removes the VRF and all VRF options, such as **authentication** , **default-cost** , **nssa** , **range** , **stub** , **virtual-link** , and **interface.**

To avoid possibly having the router ID change under a VRF, explicitly configure the router ID using the **router-id** command.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to configure VRF vrf1 and HundredGigE interface 0/2/0/0. HundredGigE interface 0/2/0/0 is bound to VRF vrf1 automatically.

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router ospf 1
RP/0/RP0/CPU0:router(config-ospf)# vrf vrf1
RP/0/RP0/CPU0:router(config-ospf-vrf)# area area1
RP/0/RP0/CPU0:router(config-ospf-vrf-ar)# interface HundredGigE 0/2/0/0
```



OSPFv3 Commands

- address-family (OSPFv3), on page 429
- area (OSPFv3), on page 430
- authentication (OSPFv3) , on page 431
- auto-cost (OSPFv3), on page 433
- capability vrf-lite (OSPFv3), on page 434
- clear ospfv3 process, on page 435
- clear ospfv3 redistribution, on page 436
- clear ospfv3 routes, on page 437
- clear ospfv3 statistics, on page 438
- cost (OSPFv3), on page 440
- database-filter all out (OSPFv3), on page 442
- dead-interval (OSPFv3), on page 443
- default-cost (OSPFv3), on page 445
- default-information originate (OSPFv3), on page 446
- default-metric (OSPFv3), on page 448
- demand-circuit (OSPFv3), on page 449
- distance ospfv3, on page 451
- distribute-list prefix-list in, on page 452
- distribute-list prefix-list out, on page 454
- domain-id (OSPFv3), on page 456
- encryption, on page 457
- flood-reduction (OSPFv3), on page 459
- graceful-restart (OSPFv3), on page 460
- hello-interval (OSPFv3), on page 461
- instance (OSPFv3), on page 463
- interface (OSPFv3), on page 465
- log adjacency changes (OSPFv3), on page 467
- maximum interfaces (OSPFv3), on page 468
- maximum paths (OSPFv3), on page 469
- maximum redistributed-prefixes (OSPFv3), on page 470
- mtu-ignore (OSPFv3), on page 472
- neighbor (OSPFv3), on page 474
- network (OSPFv3), on page 476

- nssa (OSPFv3), on page 478
- nsr (OSPFv3), on page 480
- ospfv3 name-lookup, on page 481
- packet-size (OSPFv3), on page 482
- passive (OSPFv3), on page 483
- priority (OSPFv3), on page 485
- protocol shutdown (OSPFv3), on page 486
- range (OSPFv3), on page 487
- redistribute (OSPFv3), on page 489
- retransmit-interval (OSPFv3), on page 493
- router-id (OSPFv3), on page 495
- router ospfv3, on page 497
- sham-link (OSPFv3), on page 498
- show ospfv3, on page 499
- show ospfv3 border-routers, on page 505
- show ospfv3 database, on page 507
- show ospfv3 flood-list, on page 519
- show ospfv3 interface, on page 521
- show ospfv3 message-queue, on page 524
- show ospfv3 neighbor, on page 526
- show ospfv3 request-list, on page 532
- show ospfv3 retransmission-list, on page 534
- show ospfv3 routes, on page 536
- show ospfv3 statistics rib-thread, on page 538
- show ospfv3 summary-prefix, on page 540
- show ospfv3 virtual-links, on page 542
- show protocols (OSPFv3), on page 544
- snmp context (OSPFv3), on page 546
- snmp trap (OSPFv3), on page 548
- snmp trap rate-limit (OSPFv3), on page 549
- spf prefix-priority (OSPFv3), on page 550
- stub (OSPFv3), on page 551
- stub-router, on page 552
- summary-prefix (OSPFv3), on page 554
- timers lsa arrival, on page 556
- timers pacing flood, on page 557
- timers pacing lsa-group, on page 558
- timers pacing retransmission , on page 560
- timers throttle lsa all (OSPFv3), on page 561
- timers throttle spf (OSPFv3), on page 563
- trace (OSPFv3), on page 565
- transmit-delay (OSPFv3), on page 567
- virtual-link (OSPFv3), on page 568
- vrf (OSPFv3), on page 569

address-family (OSPFv3)

To enter address family configuration mode for Open Shortest Path First Version 3 (OSPFv3), use the **address-family** command in the router ospv3 configuration mode. To disable address family configuration mode, use the **no** form of this command.

address-family ipv6 [**unicast**]

Syntax Description

ipv6	Specifies IP Version 6 (IPv6) address prefixes.
unicast	(Optional) Specifies unicast address prefixes.

Command Default

An address family is not specified.

Command Modes

Router ospfv3 configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to configure the OSPFv3 router process with IPv6 unicast address prefixes:

```
RP/0/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*

Syntax Description

area-id	Identifier of an OSPFv3 area. The *area-id* argument can be specified as either a decimal value or as an IPv4 address.

Command Default

No OSPFv3 areas are defined.

Command Modes

Router OSPFv3 configuration

OSPFv3 VRF configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

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 HundredGigE 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 HundredGigE 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** }

Syntax Description

ipsec	Specifies IP Security (IPSec). IPSec supported only for OSPFv3.
spi *spi-value*	Specifies a security policy index (SPI) value. Range is 256 to 4294967295.
md5	Enables Message Digest 5 (MD5) authentication.
sha1	Enables SHA1 authentication.
clear	(Optional) Specifies that the key be unencrypted.
password	(Optional) Specifies that the key be encrypted using a two-way algorithm.
password	Any contiguous string that can be entered from the keyboard.
disable	Disables authentication for OSPFv3 packets.

Command Default

If this command is not specified in interface configuration mode, then the interface adopts the authentication parameter specified by the area.

If this command is not specified in area configuration mode, then the interface adopts the authentication parameter specified for the process.

If this command is not specified at any level, then the interface does not use authentication.

Command Modes

Interface configuration

Area configuration

Router configuration

Virtual-link configuration

OSPFv3 VRF configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

Use the **authentication** command to specify an authentication type for the interface, which overrides the authentication specified for the area to which this interface belongs. If this command is not included in the

configuration file, the authentication configured in the area to which the interface belongs is assumed (as specified by the area **authentication** command).

The authentication type and password must be the same for all OSPFv3 interfaces that are to communicate with each other through OSPFv3.



Note

- IPSec is supported only for Open Shortest Path First version 3 (OSPFv3).
- 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**]

Syntax Description

reference-bandwidth *mbps*	(Optional) Sets the rate in Mbps (bandwidth). Range is 1 to 4294967.
disable	(Optional) Sets the link cost based only on the interface type.

Command Default

mbps : 100 Mbps

Command Modes

Router ospfv3 configuration

OSPFv3 VRF configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

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
```

capability vrf-lite (OSPFv3)

To ignore DN bit in LSAs received from peers in the given VRF and to disable automatic ABR status in that VRF, use the **capability vrf-lite** command in OSPFv3 VRF configuration mode. To disable ignoring the DN bit in LSAs and to re-enable automatic ABR status in the VRF, use the **no** form of this command.

capability vrf-lite

This command has no keywords or arguments.

Command Default

Disabled.

Command Modes

OSPFv3 VRF configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

Use the **capability vrf-lite** command when routers (sometimes called multi-vrf CE routers) are directly connected through interfaces associated with a VRF, but without being connected to other PEs through the MPLS/VPN BGP Backbone.

When OSPFv3 is enabled in a VRF, the router is always an ABR. With the capability vrf-lite command enabled, the router becomes an ABR only if it is connected to area 0 (backbone area), and there are other (non-backbone) areas enabled on this router in the given VRF.



Note Routes may be re-introduced to the VPN backbone when this command is used.

Task ID

Task ID	Operation
ospf	read, write

This example shows how to enable VRF-lite capability for OSPFv3 instance *1* under VRF *vrf1*:

```
RP/0/RP0/CPU0:router#configure
RP/0/RP0/CPU0:router(config)#router ospfv3 1
RP/0/RP0/CPU0:router(config-ospfv3)#vrf vrf1
RP/0/RP0/CPU0:router(config-ospfv3-vrf)#capability vrf-lite
```

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*] [**vrf** *vrf-name*] **process**

Syntax Description

process-name	(Optional) Name that uniquely identifies an OSPFv3 routing process. The process name is defined by the **router ospfv3** command. If this argument is included, only the specified routing process is affected. Otherwise, all OSPFv3 processes are reset.
vrf	(Optional) Specifies VPN routing and forwarding (VRF).
vrf-name	Name of a VRF.

Command Default

No default behavior or value

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

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 495 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
```

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*] [**vrf** *vrf-name*] **redistribution**

Syntax Description

process-name	(Optional) Name that uniquely identifies an OSPFv3 routing process. The process name is defined by the **router ospfv3** command. If this argument is included, only the specified routing process is affected. Otherwise, all OSPFv3 processes are reset.
vrf	(Optional) Specifies VPN routing and forwarding (VRF).
vrf-name	Name of a VRF.

Command Default

No default behavior or value

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

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*] [**vrf** *vrf-name*] **routes**

Syntax Description

process-name	(Optional) Name that uniquely identifies an OSPFv3 routing process. The process name is defined by the **router ospfv3** command. If this argument is included, only the specified routing process is affected. Otherwise, all OSPFv3 processes are reset.
vrf	(Optional) Specifies VPN routing and forwarding (VRF).
vrf-name	Name of a VRF.

Command Default

No default behavior or value

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

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*] [**vrf** *vrf-name*] **statistics** [**neighbor** [*type interface-path-id*] [*router-id*]]

Syntax Description

process-name	(Optional) Name that uniquely identifies an OSPFv3 routing process. The process name is defined by the **router ospfv3** command. If this argument is included, only the specified routing process is affected.
neighbor	(Optional) Clears counters for the specified neighbor only.
type	Interface type. For more information, use the question mark (?) online help function.
interface-path-id	Physical interface or virtual interface. **Note** Use the **show interfaces** command to see a list of all interfaces currently configured on the router. For more information about the syntax for the router, use the question mark (?) online help function.
router-id	(Optional) Specified router ID. This argument must be in 32-bit dotted-decimal notation, similar to an IPv4 address. This argument clears the counters of the specified neighbor only.
vrf	(Optional) Specifies VPN routing and forwarding (VRF).
vrf-name	Name of a VRF.

Command Default

No default behavior or value

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

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 HundredGigE interface 0/2/0/0:

```
RP/0/RP0/CPU0:router# clear ospfv3 statistics neighbor HundredGigE 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*

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

OSPFv3 VRF configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

The link-state metric is advertised as the link cost in the router link advertisement.

In general, the path cost is calculated using the following formula:

10^8 / bandwidth

Using this formula, the default path cost is 1 for OC-3 Packet-over-SONET/SDH (POS). 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 HundredGigE 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 HundredGigE 0/1/0/1
RP/0/RP0/CPU0:router(config-ospfv3-ar-if)# cost 65
```

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

Syntax Description

This command has no keywords or arguments.

Command Default

If this command is not specified in interface configuration mode, then the interface adopts the database filter parameter specified for the area.

If this command is not specified in area configuration mode, then the interface adopts the database filter parameter specified for the process.

If this command is not specified in router ospfv3 configuration mode, then the database filter is disabled and all outgoing LSAs are flooded to the interface.

Command Modes

Interface configuration

Area configuration

Router OSPFv3 configuration

OSPFv3 VRF configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

Use the **database-filter all out** command to perform the same function that the **neighbor** command (with the **database-filter** keyword) performs on a neighbor basis.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to prevent flooding of OSPFv3 LSAs to neighbors reachable through HundredGigE 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 HundredGigE 0/2/0/3
RP/0/RP0/CPU0:router(config-ospfv3-ar-if)# database-filter all out
```

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*

Syntax Description

seconds	Unsigned integer that specifies the interval (in seconds). The value must be the same for all nodes on the same network link. Range is 1 to 65535.

Command Default

If this command is not specified in interface configuration mode, then the interface adopts the dead interval parameter specified for the area.

If this command is not specified in area configuration mode, then the interface adopts the dead interval parameter specified for the process.

If this command is not specified in router ospfv3 configuration mode, then the dead interval is four times the interval set by the **hello-interval (OSPFv3)** command.

Command Modes

Interface configuration

Area configuration

Router OSPFv3 configuration

Virtual-link configuration

OSPFv3 VRF configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

Two Open Shortest Path First Version 3 (OSPFv3) routers do not become adjacent if their dead interval values differ.

If the hello interval is configured, the dead interval value must be larger than the hello interval value. The dead interval value is usually configured four times larger than the hello interval value.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to set the OSPFv3 dead interval on HundredGigE 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 HundredGigE 0/2/0/3
RP/0/RP0/CPU0:router(config-ospfv3-ar-if)# dead-interval 40
```

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*

Syntax Description

cost	Cost for the default summary route used for a stub or NSSA area. The acceptable value is a 24-bit number ranging from 1 to 16777214.

Command Default

cost : 1

Command Modes

Area configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

Use the **default-cost** command only on an Area Border Router (ABR) attached to a stub or an NSSA area.

In all routers and access servers attached to the stub area, the area should be configured as a stub area using the **stub (OSPFv3)** command in the area configuration submode. Use the **default-cost** command only on an ABR attached to the stub area. The **default-cost** command provides the metric for the summary default route generated by the ABR into the stub area.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to assign a cost of 20 to the default route sent into area 10.15.0.0:

```
RP/0/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 HundredGigE 0/3/0/1
```

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*]

Syntax Description

route-policy *policy-name*	(Optional) Specifies the route policy to apply to default information origination.
always	(Optional) Always advertises the default route regardless of whether the software has a default route.
metric *metric-value*	(Optional) Specifies a metric used for generating the default route. The default metric value is 1. The value used is specific to the protocol.
metric-type *type-value*	(Optional) Specifies an external link type associated with the default route advertised into the OSPFv3 routing domain. It can be one of the following values: **1**—Type 1 external route **2**—Type 2 external route
tag *tag-value*	(Optional) 32-bit dotted-decimal value attached to each external route. This is not used by the OSPFv3 protocol itself. It may be used to communicate information between autonomous system boundary routers (ASBRs). If a tag is not specified, then zero (0) is used.

Command Default

A default external route into an OSPFv3 routing domain is not generated.

metric-value : 1

type-value : Type 2

tag-value: 0

Command Modes

Router ospfv3 configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

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 8000 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 8000 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
```

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*

Syntax Description

value	Default metric value appropriate for the specified routing protocol.

Command Default

Built-in, automatic metric translations, as appropriate for each routing protocol

Command Modes

Router OSPFv3 configuration

OSPFv3 VRF configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

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_isp
```

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**]

Syntax Description

disable	(Optional) Disables the demand circuit configuration that may have been specified at a higher level in the configuration.

Command Default

If this command is not specified in interface configuration mode, then the interface adopts the demand circuit parameter specified for the area.

If this command is not specified in area configuration mode, then the interface adopts the demand circuit parameter specified for the process.

If this command is not specified at any level, then the interface is not a demand circuit.

Command Modes

Interface configuration

Area configuration

Router OSPFv3 configuration

Virtual-link configuration

OSPFv3 VRF configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

On point-to-point interfaces, only one end of the demand circuit must be configured with the **demand-circuit** command. Periodic hello messages are suppressed and periodic refreshes of link-state advertisements (LSAs) do not flood the demand circuit. This command allows the underlying data link layer to be closed when the topology is stable. In point-to-multipoint topology, only the multipoint end must be configured with this command.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to configure HundredGigE 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 HundredGigE 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*

Syntax Description

intra-area | **inter-area** | **external**	Type of area. It can be one of the following values: **intra-area** —All routes within an area. **inter-area** —All routes from one area to another area. **external** —All routes from other routing domains, learned by redistribution.
distance	The route administrative distance.

Command Default

distance : 110

Command Modes

Router ospfv3 configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

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
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
```

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**

Syntax Description

prefix-list-name	IP Version 6 (IPv6) prefix list name. The list defines which IPv6 prefixes are installed in the RIB.

Command Default

All routes learned by OSPFv3 are installed in the RIB.

Command Modes

Interface configuration

Router OSPFv3 configuration

OSPFv3 VRF configuration

Usage Guidelines

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 HundredGigE 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 HundredGigE 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*]]

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, isis**, **ospfv3, static**, and **connected**. The **static** keyword is used to redistribute IPv6 static routes. The **connected** keyword refers to routes that are established automatically because IPv6 is enabled on an interface. For routing protocols such as OSPFv3 and Intermediate System-to-Intermediate System (IS-IS), these routes are redistributed as external to the autonomous system.
process-id	(Optional) For the **bgp** keyword, an autonomous system number has the following ranges: • Range for 2-byte Autonomous system numbers (ASNs) is 1 to 65535. • Range for 4-byte Autonomous system numbers (ASNs) in asplain format is 1 to 4294967295. • Range for 4-byte Autonomous system numbers (ASNs) is asdot format is 1.0 to 65535.65535. For the **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 489 command are redistributed into OSPFv3.

Command Modes

Router OSPFv3 configuration

OSPFv3 VRF configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

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.

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-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 p2
RP/0/RP0/CPU0:router(config-ipv6-pfx)# deny 2064::/16 le 128
RP/0/RP0/CPU0:router(config-ipv6-pfx)# permit ::/0 le 128
!
RP/0/RP0/CPU0:router(config)# ipv6 prefix-list p3
RP/0/RP0/CPU0:router(config-ipv6-pfx)# 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
```

domain-id (OSPFv3)

To specify the Open Shortest Path First Version 3 (OSPFv3) VPN routing and forwarding (VRF) domain ID, use the **domain-id** command in VRF configuration mode. To remove the OSPFv3 VRF domain ID, use the **no** form of this command.

domain-id [**secondary**] **type** [**0005** | **0105** | **0205**] **value** *domain-id_value*

Syntax Description

secondary	(Optional) OSPFv3 secondary domain ID.
type	Primary OSPFv3 domain ID in hexadecimal format. • 0005 —Type 0x0005 • 0105 —Type 0x0105 • 0205—Type 0x0205
value	OSPF domain ID value in hexadecimal format.
domain-id-value	OSPF domain ID extended community value as a 6 byte hexadecimal number.

Command Default

No domain ID is specified.

Command Modes

VRF configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

If no value is specified for the domain ID, the default is Null (all zeros) primary domain ID. One or more secondary domain IDs can be specified.

Task ID

Task ID	Operation
ospf	read, write

This example shows how to specify a domain ID with type *0105* and value *AABBCCDDEEFF*:

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router ospf 1
RP/0/RP0/CPU0:router(config-ospf)# vrf vrf_1
RP/0/RP0/CPU0:router(config-ospf-vrf)# domain-id type 0105 value AABBCCDDEEFF
```

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*] }

Syntax Description

disable	Disables OSPFv3 packet encryption.
ipsec spi	Specifies IPSec ESP encryption and authentication with the Security Parameter Index (SPI) value. IPSec supported only for OSPFv3.
spi-value	SPI value. Range is 256 to 4294967295.
esp	Specifies Encryption Service Payload (ESP) encryption parameters.
3des	Specifies the triple DES algorithm.
aes	Specifies the Advanced Encryption Standard (AES) algorithm.
192	(Optional) Specifies the 192-bit AES algorithm
256	(Optional) Specifies the 256-bit AES algorithm
des	Specifies the Data Encryption Standard (DES) algorithm.
null	Specifies no AES algorithm.
md5	Enables Message Digest 5 (MD5) authentication.
sha1	Enables SHA1 authentication.
clear	Specifies that the key be unencrypted.
password	Specifies that the key be encrypted using a two-way algorithm.
encrypt-password	Any contiguous string that can be entered from the keyboard as the encryption password.
auth-password	Any contiguous string that can be entered from the keyboard as the authentication password.

Command Default

No default behavior or values.

Command Modes

Interface configuration

Router OSPFv3 configuration

OSPFv3 VRF configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

Use the **encryption** command to encrypt and authenticate OSPFv3 packets.

Note

- IPSec is supported only for Open Shortest Path First version 3 (OSPFv3).
- If OSPFv3 is configured along with IPsec authentication, then it is likely that adjacencies may flap on a Route Processor Fail Over (RPFO) even when NSR and/or Graceful Restart is enabled.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to encrypt and authenticate OSPFv3 packets:

```
RP/0/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**]

Syntax Description

disable	(Optional) Turns off this functionality at a specific level.

Note The **disable** keyword is not available in router ospfv3 configuration mode.

Command Default

If this command is not specified in interface configuration mode, then the interface adopts the flood reduction parameter specified by area.

If this command is not specified in area configuration mode, then the interface adopts the flood reduction parameter specified for the process.

If this command is not specified at any level, then flood reduction is disabled.

Command Modes

Interface configuration

Area configuration

Router OSPFv3 configuration

OSPFv3 VRF configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

All routers that support Open Shortest Path First Version 3 (OSPFv3) demand circuits are compatible with routers supporting flooding reduction.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to reduce the flooding of unnecessary LSAs for area 0:

```
RP/0/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 HundredGigE 0/1/0/3
RP/0/RP0/CPU0:router(config-ospfv3-ar-if)# flood-reduction
```

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*]

Syntax Description

helper disable	(Optional) Disables the routers helper support level.
interval *interval*	(Optional) Specifies the minimum interval between graceful restarts. Range is 90 to 3600 seconds.
lifetime *lifetime*	(Optional) Specifies the maximum route lifetime following a restart. Range is 90 to 3600 seconds.

Command Default

No default behavior or values.

Command Modes

Router OSPFv3 configuration

OSPFv3 VRF configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

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
```

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*

Syntax Description

seconds	Interval (in seconds). The value must be the same for all nodes on a specific network.

Command Default

If this command is not specified in interface configuration mode, then the interface adopts the hello interval parameter specified by area.

If this command is not specified in area configuration mode, then the interface adopts the hello interval parameter specified for the process.

If this command is not specified at any level, then the hello interval is 10 seconds (broadcast) or 30 seconds (non-broadcast).

Command Modes

Interface configuration

Area configuration

Router OSPFv3 configuration

Virtual-link configuration

OSPFv3 VRF configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

The hello interval value is advertised in the hello packets. The shorter the hello interval, the faster topological changes are detected, but more routing traffic ensues. This value must be the same for all routers and access servers on a specific network.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to set the interval between hello packets to 15 seconds on GigabitEthernet interface 0/3/0/2:

```
RP/0/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
```

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*

Syntax Description

instance-id	Instance identifier sent in OSPFv3 packets. Range is 0 to 255. The same value must be used by all the communicating OSPFv3 routers on a link.

Command Default

If this command is not specified in interface configuration mode, then the interface adopts the instance parameter specified by the area.

If this command is not specified in area configuration mode, then the interface adopts the instance parameter specified for the process.

If this command is not specified at any level, then the instance is 0.

Command Modes

Interface configuration

Area configuration

Router OSPFv3 configuration

OSPFv3 VRF configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

The OSPFv3 routing protocol allows multiple unrelated OSPFv3 processes to share a link by using an 8-bit "instance" value to demultiplex the protocol packets. Each OSPFv3 process sets its configured instance value in the OSPFv3 packets that it sends, and ignores received packets with instance values from other OSPFv3 processes.

Note The *instance-id* argument should not be confused with the *process-name* argument that is specified by the **router ospfv3** command. The former is an 8-bit integer that is sent to other routers as part of the OSPFv3 protocol, and the latter is a 1- to 40-character ASCII string that is significant only within a given router. The instance ID value is also unrelated to the router ID that is specified by the **router-id** command, which is a 32-bit integer value that uniquely identifies a router within an OSPFv3 routing domain.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to set the instance value for HundredGigE 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 HundredGigE 0/3/0/1
RP/0/RP0/CPU0:router(config-ospfv3-ar-if)# instance 42
```

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*

Syntax Description

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

Command Default

An interface is not defined.

Command Modes

Area configuration

OSPFv3 VRF configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

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 HundredGigE interface 0/3/0/1 is set at 40; the cost value for HundredGigE 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 HundredGigE 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 HundredGigE 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

detail	(Optional) Provides all (DOWN, INIT, 2WAY, EXSTART, EXCHANGE, LOADING, FULL) adjacency state changes.
disable	(Optional) Disables the neighbor state change messages.

Command Default

Neighbor state change messages are enabled.

Command Modes

Router OSPFv3 configuration

OSPFv3 VRF configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

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*

Syntax Description

number-interfaces	Maximum number of interfaces that can be configured for this OSPFv3 process. Range is 1 to 1024.Range is 1 to 4294967295.

Command Default

If the command is not specified, the default is 255.

If the command is not specified, the default is 1024.

Command Modes

Router OSPFv3 configuration

OSPFv3 VRF configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

Task ID

Task ID	Operations
ospf	read, write

Examples

This example shows how to allow a maximum of 500 interfaces in an OSPFv3 process:

```
RP/0/RP0/CPU0:router(config)# router ospfv3 1
RP/0/RP0/CPU0:router(config-ospfv3)# maximum interfaces 500
```

This example shows how to allow a maximum of 1500 interfaces in an OSPFv3 process:

```
RP/0/RP0/CPU0:router(config)# router ospfv3 1
RP/0/RP0/CPU0:router(config-ospfv3)# maximum interfaces 1500
```

maximum paths (OSPFv3)

To control the maximum number of parallel routes that the Open Shortest Path First Version 3 (OSPFv3) can support, use the **maximum paths** command in an appropriate configuration mode. To remove the **maximum paths** command from the configuration file and restore the system to its default condition with respect to the routing protocol, use the **no** form of this command.

maximum paths *maximum-routes-number*

Syntax Description

maximum-routes-number	Maximum number of parallel routes that OSPFv3 can install in a routing table. Range is 1 to 16. **Note** The maximum number of paths that can be configured is 16.

Command Default

16 paths

Command Modes

Router OSPFv3 configuration

VRF configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

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**]

Syntax Description

limit	Maximum number of IP Version 6 (IPv6) prefixes that are allowed to be redistributed into OSPFv3, or, if the **warning-only** keyword is present, sets the number of prefixes allowed to be redistributed into OSPFv3 before the system logs a warning message. Range is 1 to 4294967295. **Note** If the **warning-only** keyword is also configured, this value does not limit redistribution; it is simply the number of redistributed prefixes that, when reached, causes a warning message to be logged.
threshold	(Optional) Percentage of the value set for the maximum number of redistributed prefixes that, when reached, causes a warning message to be logged.
warning-only	(Optional) Causes a warning to be logged when the number of routes defined by the *limit* argument have been redistributed. Additional redistribution is not prevented.

Command Default

limit : 10240

threshold : 75 percent

Command Modes

Router OSPFv3 configuration

OSPFv3 VRF configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

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**]

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

OSPFv3 VRF configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

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 HundredGigE 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 HundredGigE 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**]

Syntax Description

ipv6-address	Link- local IP Version 6 (IPv6) address of the neighbor. This argument must be in the form documented in RFC 2373, in which the address is specified in hexadecimal using 16-bit values between colons.
priority *number*	(Optional) Specifies an 8-bit number indicating the router priority value of the nonbroadcast neighbor associated with the IP address specified. The **priority** keyword does not apply to point-to-multipoint interfaces.
poll-interval *seconds*	(Optional) Specifies an unsigned integer value (in seconds) reflecting the poll interval. RFC 1247 recommends that this value be much larger than the hello interval. The **poll-interval** keyword does not apply to point-to-multipoint interfaces.
cost *number*	(Optional) Assigns a cost to the neighbor, in the form of an integer from 1 to 65535. Neighbors with no specific cost configured assume the cost of the interface, based on the **cost** command. On point-to-multipoint interfaces, **cost** *number* is the only keyword and argument combination that works. The **cost** keyword does not apply to nonbroadcast multiaccess (NBMA) networks.
database-filter all out	(Optional) Filters outgoing link-state advertisements (LSAs) to an OSPFv3 neighbor.

Command Default

No configuration is specified.

priority *number* : 0

poll-interval *seconds* : 120 seconds (2 minutes)

Command Modes

Interface configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

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 HundredGigE 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
```

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** } }

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

OSPFv3 VRF configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

Use the **network** command to configure broadcast networks as NBMA networks when, for example, routers in your network do not support multicast addressing.

Most times, it is assumed that when you configure NBMA networks as either broadcast or nonbroadcast, there are virtual circuits from every router to every router or fully meshed network. However, there are other configurations where this assumption is not true; for example, a partially meshed network. In these cases, you can configure the OSPFv3 network type as a point-to-multipoint network. Routing between two routers that are not directly connected goes through the router that has virtual circuits to both routers. You need not configure neighbors when using this command.

If the **network** command is issued on an interface that does not allow it, this command is ignored.

OSPFv3 has two features related to point-to-multipoint networks. One feature applies to broadcast networks and the other feature applies to nonbroadcast networks:

- On point-to-multipoint, broadcast networks, you can use the **neighbor** command, and you must specify a cost to that neighbor.
- On point-to-multipoint, nonbroadcast networks, you must use the **neighbor** command to identify neighbors. Assigning a cost to a neighbor is optional.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to configure an Ethernet interface as point-to-point:

```
RP/0/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 HundredGigE 0/1/0/3
RP/0/RP0/CPU0:router(config-ospfv3-ar-if)# network point-to-point
```

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**]

Syntax Description

no-redistribution	(Optional) Imports routes into the normal areas, but not into the NSSA area, by the **redistribute** command when the router is an NSSA area border router (ABR).
default-information-originate	(Optional) Generates a Type 7 default into the NSSA area. This keyword takes effect only on an NSSA ABR or NSSA autonomous system boundary router (ASBR).
metric *metric-value*	(Optional) Specifies a metric used for generating the default route. If you do not specify a default route metric value using the **nssa** and **defaultmetric** commands, the default metric value is 10. The value used is specific to the protocol.
metric-type *type-value*	(Optional) Specifies an external link type associated with the default route advertised into the Open Shortest Path First Version 3 (OSPFv3) routing domain. It can be one of the following values: **1**—Type 1 external route **2**—Type 2 external route
no-summary	(Optional) Prevents an (ABR) from sending summary link advertisements into the NSSA area.

Command Default

No NSSA area is defined.

If you do not specify a value using the **default-metric** command, the default metric value is 10.

The default *type-value* is Type 2 external route.

Command Modes

Area configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

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**]

Syntax Description

This command has no keywords or arguments.

Command Default

NSR is enabled.

Command Modes

OSPFv3 Router configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

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 497	
show ospfv3, on page 499	

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

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 mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

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*

Syntax Description

bytes	Size in bytes. Range is 256 to 10000 bytes.

Command Default

If not specified, the default packet size is 1500 bytes.

Command Modes

Router OSPFv3 configuration

Area configuration

Interface configuration

OSPFv3 VRF configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

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**]

Syntax Description

disable	(Optional) Sends OSPFv3 updates.

Note
The **disable** keyword is not available in router ospfv3 configuration mode.

Command Default

If this command is not specified in interface configuration mode, then the interface adopts the passive parameter specified by the area.

If this command is not specified in area configuration mode, then the interface adopts the passive parameter specified for the process.

If this command is not specified at any level, then the passive parameter is disabled and OSPFv3 updates are sent on the interface.

Command Modes

Interface configuration

Area configuration

Router OSPFv3 configuration

OSPFv3 VRF configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

OSPF routing information is neither sent nor received through the specified router interface. The specified interface address appears as a stub network in the OSPF domain.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows that OSPFv3 updates run over HundredGigE interface 0/3/0/0, 0/2/0/0, and 0/2/0/2. All other interfaces suppress sending OSPFv3 updates because they are in passive mode.

```
RP/0/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 HundredGigE 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 HundredGigE 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 HundredGigE 0/2/0/0
RP/0/RP0/CPU0:router(config-ospfv3-ar-if)# exit
RP/0/RP0/CPU0:router(config-ospfv3-ar)# interface HundredGigE 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 HundredGigE 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*

Syntax Description

value	8-bit unsigned integer indicating the router priority value. Range is 0 to 255.

Command Default

If this command is not specified in interface configuration mode, then the interface adopts the priority parameter specified by the area.

If this command is not specified in area configuration mode, then the interface adopts the priority parameter specified by the process.

If this command is not specified at any level, then the default priority is 1.

Command Modes

Interface configuration

Area configuration

Router OSPFv3 configuration

OSPFv3 VRF configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

When two routers attached to a network both attempt to become the designated router, the one with the higher router priority takes precedence. If there is a tie, the router with the higher router ID takes precedence. A router with a router priority set to zero is ineligible to become the designated router or backup designated router. Router priority is configured only for interfaces to broadcast and nonbroadcast multiaccess (NBMA) networks.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to set the router priority value to 4 on HundredGigE interface 0/1/0/1:

```
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 HundredGigE 0/1/0/1
RP/0/RP0/CPU0:router(config-ospfv3-ar-if)# priority 4
```

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

Command Default

None

Command Modes

ospfv3 configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

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*]

Syntax Description

ipv6-prefix	Summary prefix designated for a range of IP Version 6 (IPv6) prefixes. This argument must be in the form documented in RFC 2373, in which the address is specified in hexadecimal using 16-bit values between colons.
/ *prefix-length*	Length of the IPv6 prefix. A decimal value that indicates how many of the high-order contiguous bits of the address compose the prefix (the network portion of the address). A slash must precede the decimal value.
advertise	(Optional) Sets the address range status to advertise and generates a Type 3 summary link-state advertisement (LSA).
not-advertise	(Optional) Sets the address range status to DoNotAdvertise. The Type 3 summary LSA is suppressed and the component networks remain hidden from other networks.
cost *number*	(Optional) Specifies a cost for the range. Range is 1 to 16777214.

Command Default

Routes are not consolidated and summarized for an area.

Command Modes

Area configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

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*]

Local Interface Routes
redistribute connected [**metric** *metric-value*] [**metric-type** { **1** | **2** }] [**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*]

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*]

Static
redistribute static [**metric** *metric-value*] [**metric-type** { **1** | **2** }] [**policy** *policy-name*] [**tag** *tag-value*]

Syntax Description

bgp	Distributes routes from the BGP protocol.
process-id	For the **bgp** keyword, an autonomous system number has the following ranges: • Range for 2-byte Autonomous system numbers (ASNs) is 1 to 65535. • Range for 4-byte Autonomous system numbers (ASNs) in asplain format is 1 to 4294967295. • Range for 4-byte Autonomous system numbers (ASNs) is asdot format is 1.0 to 65535.65535. For the **isis** keyword, an IS-IS instance name from which routes are to be redistributed. The value takes the form of a string. A decimal number can be entered, but it is stored internally as a string. For the **ospf** keyword, an OSPF instance name from which routes are to be redistributed. The value takes the form of a string. A decimal number can be entered, but it is stored internally as a string.
metric *metric-value*	(Optional) Specifies the metric used for the redistributed route. Range is 1 to 16777214. Use a value consistent with the destination protocol.

metric-type { **1** | **2** }	(Optional) Specifies the external link type associated with the route advertised into the OSPF routing domain. It can be one of two values: • **1**—Type 1 external route • **2**—Type 2 external route If no **metric-type** is specified, the default is Type 2 external routes.
tag *t ag-value*	(Optional) Specifies the 32-bit dotted-decimal value attached to each external route. This value is not used by the OSPF protocol itself, but is carried in the External LSAs. Range is 0 to 4294967295.
policy *policy-name*	(Optional) Specifies the identifier of a configured policy. A policy is used to filter the importation of routes from this source routing protocol to OSPF.
connected	Distributes routes that are established automatically by virtue of having enabled IP on an interface.
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

OSPFv3 VRF configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

Note When redistributing routes (into OSPF) using both command keywords for setting or matching of attributes and a route policy, the routes are run through the route policy first, followed by the keyword matching and setting.

Disabling or changing the arguments of any keyword does not affect the state of other keywords.

In general, route redistribution from Level 1 to Level 2 is automatic. You might want to use this command to better control which Level 1 routes can be redistributed into Level 2.

The redistibution of Level 2 routes into Level 1 is called *route leaking*. Route leaking is disabled by default. That is, Level 2 routes are not automatically included in Level 1 link-state protocols. If you want to leak Level 2 routes into Level 1, you must enable that behavior by using this command.

Redistribution from Level 1 into Level 1 and from Level 2 into Level 2 is not allowed.

A router receiving a link-state packet with an internal metric considers the cost of the route from itself to the redistributing router plus the advertised cost to reach the destination. An external metric considers only the advertised metric to reach the destination.

Redistributed routing information should always be filtered by the **distribute-list prefix-list out** command. Use of this command ensures that only those routes intended by the administrator are passed along to the receiving routing protocol.

OSPFv3 Considerations

Whenever you use the **redistribute** or the **default-information** command to redistribute routes into an OSPFv3 routing domain, the router automatically becomes an ASBR. However, an ASBR does not, by default, generate a default route into the OSPFv3 routing domain.

When routes are redistributed between OSPFv3 processes, no OSPFv3 metrics are preserved.

When routes are redistributed into OSPF and no metric is specified with the **metric** keyword, OSPF uses 20 as the default metric for routes from all protocols except BGP routes, which get a metric of 1. Furthermore, when the router redistributes from one OSPFv3 process to another OSPFv3 process on the same router, and if no default metric is specified, the metrics in one process are carried to the redistributing process.

BGP Considerations

The only connected routes affected by this command are the routes not specified by the **network** (BGP) command.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to cause static routes to be redistributed into an OSPFv3 domain:

```
RP/0/RP0/CPU0:router(config)# router ospfv3 109
RP/0/RP0/CPU0:router(config-ospfv3)# redistribute isis level-1
```

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*

Syntax Description

seconds	Time (in seconds) between retransmissions. It must be greater than the expected round-trip delay between any two routers on the attached network. Range is 1 to 65535 seconds.

Command Default

If this command is not specified in interface configuration mode, then the interface adopts the retransmit interval parameter specified by the area.

If this command is not specified in area configuration mode, then the interface adopts the retransmit interval parameter specified by the process.

If this command is not specified at any level, then the default retransmit interval is 5 seconds.

Command Modes

Interface configuration

Area configuration

Router OSPFv3 configuration

Virtual-link configuration

OSPFv3 VRF configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

When a router sends an LSA to its neighbor, it keeps the LSA until it receives the acknowledgment message. If the router receives no acknowledgment, it resends the LSA.

The setting of this parameter should be conservative, or needless retransmission results. The value should be larger for serial lines and virtual links.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to set the retransmit interval value to 8 seconds while in interface configuration mode:

```
RP/0/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 HundredGigE 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*

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

OSPFv3 VRF configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

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.

In VRF configuration mode, it is mandatory to configure a router ID manually. Otherwise, the OSPFv3 process will not become operational in the VRF.

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
```

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*

Syntax Description

process-name	Name that uniquely identifies an OSPFv3 routing process. The process name is any alphanumeric string no longer than 40 characters.

Command Default

No OSPFv3 routing process is defined.

Command Modes

XR Config mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

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
```

sham-link (OSPFv3)

To configure an Open Shortest Path First version 3 (OSPFv3) sham link between two provider edge routers on a non-default VRF, use the **sham-link** command in OSPFv3 router area sub-configuration mode. To terminate an OSPFv3 sham link, use the **no** form of this command.

sham-link *source-address destination-address*

Syntax Description

source-address	IP address of the local (source) sham-link endpoint specified in four-part, dotted-decimal notation.
destination-address	IP address of the remote (destination) sham-link endpoint specified in four-part, dotted-decimal notation.

Command Default

No sham link is configured.

Command Modes

OSPFv3 router area sub-configuration.

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

Use the **sham-link** command to configure a point-to-point connection between two provider edge (PE) routers creating an interconnect between two VPN sites (VPN backbone). Sham links are configured on PE provider edge (PE) routers in a Multiprotocol Label Switching (MPLS) VPN backbone.

Task ID

Task ID	Operations
ospf	read, write

Examples

This example shows how to configure an OSPFv3 sham link:

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router ospfv3 ospfv3
RP/0/RP0/CPU0:router(config_ospfv3)# vrf 1
RP/0/RP0/CPU0:router(config_ospfv3_vrf)# area 1
RP/0/RP0/CPU0:router(config_ospfv3_vrf_ar)# sham-link 100::1 200::1
RP/0/RP0/CPU0:router(config_ospf_vrf_ar_sl)# cost 23
```

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*] [**vrf** {**all** *vrf-name*}]**sham-links**

Syntax Description

process-name	(Optional) Name that uniquely identifies an OSPFv3 routing process. The process name is defined by the **router ospfv3** command. If this argument is included, only information for the specified routing process is displayed.
vrf	Specifies an OSPF VPN routing and forwarding (VRF) instance.
all	Display all the VRFs, excluding the default VRF.
vrf-name	Specifies the name of the of the OSPFv3 VRF.
sham-links	Displays OSPFv3 Sham-link information.

Command Default

None

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

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 msecs
Minimum hold time between two consecutive SPFs 10000 msecs
Maximum wait time between two consecutive SPFs 10000 msecs
Initial LSA throttle delay 0 msecs
Minimum hold time for LSA throttle 5000 msecs
Maximum wait time for LSA throttle 5000 msecs
```

```
Minimum LSA arrival 1000 msecs
LSA group pacing timer 240 secs
Interface flood pacing timer 33 msecs
Retransmission pacing timer 66 msecs
Maximum number of configured interfaces 255
Number of external LSA 1. Checksum Sum 0x004468
Number of areas in this router is 1. 1 normal 0 stub 0 nssa
 Area BACKBONE(0) (Inactive)
  Number of interfaces in this area is 1
  SPF algorithm executed 1 times
  Number of LSA 3. Checksum Sum 0x018109
  Number of DCbitless LSA 0
  Number of indication LSA 0
  Number of DoNotAge LSA 0
  Flood list length 0
```

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 msecs
Minimum hold time between two consecutive SPFs 10000 msecs
Maximum wait time between two consecutive SPFs 10000 msecs
Initial LSA throttle delay 0 msecs
Minimum hold time for LSA throttle 5000 msecs
Maximum wait time for LSA throttle 5000 msecs
Minimum LSA arrival 1000 msecs
LSA group pacing timer 240 secs
Interface flood pacing timer 33 msecs
Retransmission pacing timer 66 msecs
Maximum number of configured interfaces 255
Number of external LSA 1. Checksum Sum 0x004468
Number of areas in this router is 1. 1 normal 0 stub 0 nssa
 Area BACKBONE(0) (Inactive)
  Number of interfaces in this area is 1
  SPF algorithm executed 1 times
  Number of LSA 3. Checksum Sum 0x018109
  Number of DCbitless LSA 0
  Number of indication LSA 0
  Number of DoNotAge LSA 0
  Flood list length 0
```

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 msecs
Minimum hold time between two consecutive SPFs 10000 msecs
Maximum wait time between two consecutive SPFs 10000 msecs
Initial LSA throttle delay 0 msecs
Minimum hold time for LSA throttle 5000 msecs
Maximum wait time for LSA throttle 5000 msecs
Minimum LSA arrival 1000 msecs
LSA group pacing timer 240 secs
Interface flood pacing timer 33 msecs
Retransmission pacing timer 66 msecs
Maximum number of configured interfaces 255
Number of external LSA 1. Checksum Sum 0x004468
Number of areas in this router is 1. 1 normal 0 stub 0 nssa
 Area BACKBONE(0) (Inactive)
  Number of interfaces in this area is 1
  SPF algorithm executed 1 times
  Number of LSA 3. Checksum Sum 0x018109
  Number of DCbitless LSA 0
  Number of indication LSA 0
  Number of DoNotAge LSA 0
  Flood list length 0
```

This table describes the significant fields shown in the display.

Table 71: show ospfv3 Field Descriptions

Field	Description
Routing Process “ospfv3 test” with ID	OSPFv3 process name.
It is	Types are internal, area border, or autonomous system boundary.
Redistributing External Routes from	Lists of redistributed routes, by protocol.
Maximum number of redistributed prefixes	Number of redistributed prefixes.
Threshold for warning message	Warning message threshold.
Initial SPF schedule delay	Delay time of SPF calculations.
Minimum hold time between two consecutive SPFs	Minimum hold time between consecutive SPFs.
Maximum wait time between two consecutive SPFs	Maximum wait time between consecutive SPFs.
Initial LSA throttle delay	Delay time of LSA throttle.
Maximum hold time for LSA throttle	After initial throttle delay, the LSA generation is backed off by hold interval.
Maximum wait time for LSA throttle	Maximum throttle delay for LSA generation.
Minimum LSA arrival	Minimum LSA arrival.
LSA group pacing timer	Configured LSA group pacing timer (in seconds).

Field	Description
Interface flood pacing timer	Flooding pacing interval.
Retransmission pacing timer	Retransmission pacing interval.
Maximum number of configured interfaces	Maximum number of configured interfaces.
Number of external LSA	Number of external LSAs.
Number of areas in this router is	Number of areas configured for the router.
Number of interfaces in this area is	Number of interfaces in the area.
SPF algorithm executed *n* times	Times SPF algorithm was executed.
Number of LSA	Number of LSAs.
Number of DCbitless LSA	Number of DCbitless LSAs.
Number of indication LSA	Number of indication LSAs.
Number of DoNotAge LSA	Number of do-not-age LSAs.
Flood list length	Flood list length.

This is sample output from the show ospfv3 vrf command that displays domain -id configuration:

```
RP/0/RP0/CPU0:router#show ospfv3 0 vrf V1
Mon May 10 14:52:31.332 CEST

 Routing Process "ospfv3 0" with ID 100.0.0.2 VRF V1
 It is an area border and autonomous system boundary router
 Redistributing External Routes from,
    bgp 1
    Maximum number of redistributed prefixes 10240
    Threshold for warning message 75%
 Primary Domain ID:
   0x0005:0xcafe00112233
 Secondary Domain ID:
   0x0105:0xbeef00000001
   0x0205:0xbeef00000002
 Initial SPF schedule delay 5000 msecs
 Minimum hold time between two consecutive SPFs 10000 msecs
 Maximum wait time between two consecutive SPFs 10000 msecs
 Initial LSA throttle delay 0 msecs
 Minimum hold time for LSA throttle 5000 msecs
 Maximum wait time for LSA throttle 5000 msecs
 Minimum LSA arrival 1000 msecs
 LSA group pacing timer 240 secs
 Interface flood pacing timer 33 msecs
 Retransmission pacing timer 66 msecs
 Maximum number of configured interfaces 255
 Maximum number of configured paths 16
 Number of external LSA 2. Checksum Sum 0x015bb3
 Number of areas in this router is 1. 1 normal 0 stub 0 nssa
 Auto cost is enabled. Reference bandwidth 100
    Area BACKBONE(0)
        Number of interfaces in this area is 1
        SPF algorithm executed 2 times
```

```
        Number of LSA 4. Checksum Sum 0x02629d
        Number of DCbitless LSA 0
        Number of indication LSA 0
        Number of DoNotAge LSA 0
        Flood list length 0
```

This is sample output from the show ospfv3 vrf command that displays vrf-lite configuration:

```
RP/0/RP0/CPU0:router#show ospfv3 0 vrf V2
Mon May 10 18:01:38.654 CEST

 Routing Process "ospfv3 0" with ID 2.2.2.2 VRF V2
 VRF lite capability is enabled
 Initial SPF schedule delay 5000 msecs
 Minimum hold time between two consecutive SPFs 10000 msecs
 Maximum wait time between two consecutive SPFs 10000 msecs
 Initial LSA throttle delay 0 msecs
 Minimum hold time for LSA throttle 5000 msecs
 Maximum wait time for LSA throttle 5000 msecs
 Minimum LSA arrival 1000 msecs
 LSA group pacing timer 240 secs
 Interface flood pacing timer 33 msecs
 Retransmission pacing timer 66 msecs
 Maximum number of configured interfaces 255
 Maximum number of configured paths 16
 Number of external LSA 0. Checksum Sum 00000000
 Number of areas in this router is 0. 0 normal 0 stub 0 nssa
 Auto cost is enabled. Reference bandwidth 100
```

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 msecs
 Minimum hold time between two consecutive SPFs 10000 msecs
 Maximum wait time between two consecutive SPFs 10000 msecs
 Initial LSA throttle delay 0 msecs
 Minimum hold time for LSA throttle 5000 msecs
 Maximum wait time for LSA throttle 5000 msecs
 Minimum LSA arrival 1000 msecs
 LSA group pacing timer 240 secs
 Interface flood pacing timer 33 msecs
 Retransmission pacing timer 66 msecs
 Maximum number of configured interfaces 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
```

The following is a sample output from the **show ospfv3** command with **sham-links** keyword:

```
RP/0/RP0/CPU0:router# show ospfv3 vrf vrf1 sham-links

Sham Links for OSPFv3 100, VRF vrf1

Sham Link OSPF_SL1 to address 300::1 is up
Area 2, source address 100::1
IfIndex = 2
  Run as demand circuit
  DoNotAge LSA allowed., Cost of using 1
  Transmit Delay is 1 sec, State POINT_TO_POINT,
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
    Hello due in 00:00:08
    Adjacency State FULL (Hello suppressed)
    Number of DBD retrans during last exchange 0
    Index 2/2, retransmission queue length 0, number of retransmission 0
    First 0(0)/0(0) Next 0(0)/0(0)
    Last retransmission scan length is 0, maximum is 0
    Last retransmission scan time is 0 msec, maximum is 0 msec
Sham Link OSPF_SL0 to address 200::1 is up
Area 2, source address 100::1
IfIndex = 2
  Run as demand circuit
  DoNotAge LSA allowed., Cost of using 1
  Transmit Delay is 1 sec, State POINT_TO_POINT,
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
    Hello due in 00:00:01
    Adjacency State FULL (Hello suppressed)
    Number of DBD retrans during last exchange 0
    Index 3/3, retransmission queue length 0, number of retransmission 0
    First 0(0)/0(0) Next 0(0)/0(0)
    Last retransmission scan length is 0, maximum is 0
    Last retransmission scan time is 0 msec, maximum is 0 msec
```

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*] [**vrf** {**all** *vrf-name*}]**border-routers** [*router-id*]

Syntax Description

process-name	(Optional) Name that uniquely identifies an OSPFv3 routing process. The process name is defined by the **router ospfv3** command. If this argument is included, only information for the specified routing process is displayed.
vrf	Specifies an OSPF VPN routing and forwarding (VRF) instance.
all	Display all the VRFs, excluding the default VRF.
vrf-name	Specifies the name of the of the OSPFv3 VRF.
router-id	(Optional) 32-bit router ID value specified in four-part, dotted-decimal notation.

Command Default

No default behavior or values

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

Task ID

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, HundredGigE 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 72: show ospf border-routers Field Descriptions

Field	Description
i	Type of this route; i indicates an intra-area route, I an inter-area route.
10.0.0.207	Router ID of destination.
[1]	Cost of using this route.
fe80::3034:30ff:fe33:3742	Next-hop toward the destination.
HundredGigE 0/3/0/0	Packets destined for fe80::3034:30ff:fe33:3742 are sent over HundredGigE 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.

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*]] [**vrf** {**all** *vrf-name*}]**database**
show ospfv3 [*process-name* [*area-id*]] [**vrf** {**all** *vrf-name*}]**database**[**adv-router** [*router-id*]]
show ospfv3 [*process-name* [*area-id*]] [**vrf** {**all** *vrf-name*}]**database** [**database-summary**]
show ospfv3 [*process-name* [*area-id*]] [**vrf** {**all** *vrf-name*}]**database** [**external**] [*link-state-id*]
show ospfv3 [*process-name* [*area-id*]] [**vrf** {**all** *vrf-name*}]**database** [**external**] [*link-state-id*] [**internal**] [**adv-router** [*router-id*]]
show ospfv3 [*process-name* [*area-id*]] [**vrf** {**all** *vrf-name*}]**database** [**external**] [*link-state-id*] [**internal**] [**self-originate**]
show ospfv3 [*process-name* [*area-id*]] [**vrf** {**all** *vrf-name*}]**database** [**grace**] [*link-state-id*] [**adv-router** [*router-id*]] [**internal**] [**self-originate**]
show ospfv3 [*process-name* [*area-id*]][**vrf** {**all** *vrf-name*}] **database** [**inter-area prefix**] [*link-state-id*]
show ospfv3 [*process-name* [*area-id*]] **vrf** *vrf-name* **database** [**vrf** {**all** *vrf-name*}] [*link-state-id*] [**internal**] [**adv-router** [*router-id*]]
show ospfv3 [*process-name* [*area-id*]] [**vrf** {**all** *vrf-name*}]**database** [**inter-area prefix**] [*link-state-id*] [**internal**] [**self-originate**]
show ospfv3 [*process-name* [*area-id*]][**vrf** {**all** *vrf-name*}] **database** [**inter-area router**] [*link-state-id*]
show ospfv3 [*process-name* [*area-id*]] [**vrf** {**all** *vrf-name*}]**database** [**inter-area router**] [*link-state-id*] [**internal**] [**adv-router** [*router-id*]]
show ospfv3 [*process-name* [*area-id*]] [**vrf** {**all** *vrf-name*}]**database** [**inter-area router**] [*link-state-id*] [**internal**] [**self-originate**]
show ospfv3 [*process-name* [*area-id*]] [**vrf** {**all** *vrf-name*}]**database** [**link**] [*link-state-id*]
show ospfv3 [*process-name* [*area-id*]] [**vrf** {**all** *vrf-name*}]**database** [**link**] [*link-state-id*] [**internal**] [**adv-router** [*router-id*]]
show ospfv3 [*process-name* [*area-id*]] [**vrf** {**all** *vrf-name*}]**database** [**link**] [*link-state-id*] [**internal**] [**self-originate**]
show ospfv3 [*process-name* [*area-id*]] [**vrf** {**all** *vrf-name*}]**database** [**network**] [*link-state-id*]
show ospfv3 [*process-name* [*area-id*]][**vrf** {**all** *vrf-name*}] **database** [**network**] [*link-state-id*] [**internal**] [**adv-router** [*router-id*]]
show ospfv3 [*process-name* [*area-id*]] [**vrf** {**all** *vrf-name*}]**database** [**network**] [*link-state-id*] [**internal**] [**self-originate**]
show ospfv3 [*process-name* [*area-id*]][**vrf** {**all** *vrf-name*}] **database** [**nssa-external**] [*link-state-id*]
show ospfv3 [*process-name* [*area-id*]][**vrf** {**all** *vrf-name*}] **database** [**nssa-external**] [*link-state-id*] [**internal**] [**adv-router** [*router-id*]]
show ospfv3 [*process-name* [*area-id*]] [**vrf** {**all** *vrf-name*}]**database** [**nssa-external**] [*link-state-id*] [**internal**] [**self-originate**]
show ospfv3 [*process-name* [*area-id*]][**vrf** {**all** *vrf-name*}] **database** [**prefix**] [**ref-lsa**] [**router** | **network**] [*link-state-id*] [**internal**] [**adv-router** [*router-id*]]
show ospfv3 [*process-name* [*area-id*]][**vrf** {**all** *vrf-name*}] **database** [**prefix**] [**ref-lsa**] [**router** | **network**] [*link-state-id*] [**internal**] [**self-originate**]
show ospfv3 [*process-name* [*area-id*]] [**vrf** {**all** *vrf-name*}]**database** [**prefix**] [*link-state-id*]
show ospfv3 [*process-name* [*area-id*]][**vrf** {**all** *vrf-name*}] **database** [**prefix**] [*link-state-id*] [**internal**] [**adv-router** [*router-id*]]

show ospfv3 [*process-name* [*area-id*]][**vrf** {**all** *vrf-name*}] **database** [**prefix**] [*link-state-id*] [**internal**] [**self-originate**]
show ospfv3 [*process-name* [*area-id*]] [**vrf** {**all** *vrf-name*}]**database** [**router**] [*link-state-id*]
show ospfv3 [*process-name area-id*] [**vrf** {**all** *vrf-name*}]**database [router]** [**adv-router** [*router-id*]]
show ospfv3 [*process-name* [*area-id*]] [**vrf** {**all** *vrf-name*}]**database** [**router**] [*link-state-id*] [**internal**] [**self-originate**]
show ospfv3 [*process-name* [*area-id*]][**vrf** {**all** *vrf-name*}] **database** [**self-originate**]
show ospfv3 [*process-name* [*area-id*]] [**vrf** {**all** *vrf-name*}]**database** [**unknown** [**area** | **as** | **link**]] [*link-state-id*] [**internal**] [**adv-router** [*router-id*]] [**self-originate**]

Syntax Description

process-name	(Optional) Name that uniquely identifies an OSPFv3 routing process. The process name is defined by the **router ospfv3** command. If this argument is included, only information for the specified routing process is displayed.
area-id	(Optional) Area number used to define the particular area.
adv-router [*router-id*]	(Optional) Displays all link-state advertisements (LSAs) of the specified router.
asbr-summary	(Optional) Displays information only about the Autonomous System Boundary Router (ASBR) summary LSAs.
database-summary	(Optional) Displays how many of each type of LSA are in the database for each area and the total.
external	(Optional) Displays information only about external LSAs.
grace	(Optional) Displays information about the state for the graceful restart link.
internal	(Optional) Displays information only about internal LSAs.
self-originate	(Optional) Displays only self-originated LSAs (from the local router).
link-state-id	(Optional) LSA ID that uniquely identifies the LSA. For network LSAs and link LSAs, this ID is the interface ID of the link of the router originating the LSA.
inter-area prefix	(Optional) Displays information only about the interarea prefix LSAs.
inter-area router	(Optional) Displays information only about the interarea router LSAs.
link	(Optional) Displays information only about the link LSAs.
network	(Optional) Displays information only about the network LSAs.
nssa-external	(Optional) Displays information only about the not-so-stubby area (NSSA) external LSAs.
prefix	(Optional) Displays information only about the prefix LSAs.
ref-lsa	(Optional) Displays referenced LSA information.
router	(Optional) Displays information only about the router LSAs.
unknown	(Optional) Displays information only about unknown LSAs.
area	(Optional) Displays information only about the area LSAs.

as	(Optional) Displays information only about the autonomous system LSAs.
vrf	Specifies an OSPF VPN routing and forwarding (VRF) instance.
all	Display all the VRFs, excluding the default VRF.
vrf-name	Specifies the name of the of the OSPFv3 VRF.

Command Default

No default behavior or values

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

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 73: show ospfv3 database Field Descriptions

Field	Description
ADV Router	ID of advertising router.
Age	Link-state age.
Seq#	Link-state sequence number (detects old or duplicate LSAs).
Fragment ID	Router LSA fragment ID.
Link count	Number of links described.
Bits	B indicates that the router is an area border router. E indicates that the router is an autonomous system boundary router. V indicates that the router is a virtual link endpoint. W indicates that the router is a wildcard multicast receiver.
Link ID	Unique LSA ID.
Rtr count	Number of routers attached to the link.
Prefix	Prefix of the route being described.
Dest RtrID	Router ID of the router being described.
Interface	Link described by the LSA.
Ref-lstype	LSA type of the LSA being referenced.
Ref-LSID	LSA ID of the LSA being referenced.

The following is sample output from the **show ospfv3 database** command with the **external** keyword:

```
RP/0/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 74: 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 75: show ospfv3 database inter-area prefix Field Descriptions

Field	Description
OSPFv3 Router with ID	Router ID number.
Process ID	OSPFv3 process name.
LS age	Link-state age.
LS Type	Link-state type.
Link State ID	Link-state ID.
Advertising Router	ID of advertising router.
LS Seq Number	Link-state sequence (detects old or duplicate LSAs).
Checksum	Link-state checksum (Fletcher checksum of the complete contents of the LSA).
Length	Length (in bytes) of the LSA.
Metric	Link-state metric.
Prefix Address	IPv6 prefix of the route being described.
Prefix Length	IPv6 prefix length of the route being described.
Options	LA indicates that the prefix is a local address. MC indicates the prefix is multicast capable. NU indicates that the prefix is not unicast capable. P indicates that the prefix should be propagated at a not-so-stubby area (NSSA) area border.

The following is sample output from the **show ospfv3 database** command with the **inter-area router** keyword:

```
RP/0/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 76: show ospfv3 database inter-area routerField Descriptions

Field	Description
OSPFv3 Router with ID	Router ID number.
Process ID	OSPFv3 process name.
LS age	Link-state age.
Options	Type of service options (Type 0 only): DC—Supports demand circuits.E—Capable of processing external LSAs. MC—Forwards IP multicast. N—Supports Type 7 LSAs. R—Router is active. V6—Include in IPv6 routing calculations.
LS Type	Link-state type.
Link State ID	Link-state ID.
Advertising Router	ID of the advertising router.
LS Seq Number	Link-state sequence (detects old or duplicate LSAs).
Checksum	Link-state checksum (Fletcher checksum of the complete contents of the LSA.)
Length	Length (in bytes) of the LSAs.
Metric	Link-state metric.
Destination Router ID	Router ID of the router being described.

The following is sample output from the **show ospfv3 database** command with the **link** keyword:

```
RP/0/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 77: show ospfv3 database link Field Descriptions

Field	Description
OSPFv3 Router with ID	Router ID number.
Process ID	OSPFv3 process name.
LS age	Link-state age.
Options	Type of service options (Type 0 only): DC—Supports demand circuits.E—Capable of processing external LSAs.MC—Forwards IP multicast.N—Supports type-7 LSAs.R—Router is active.V6—Include in IPv6 routing calculations.
LS Type	Link-state type.
Link State ID	Link-state ID (Interface ID).
Advertising Router	ID of the advertising router.
LS Seq Number	Link-state sequence (detects old or duplicate LSAs).
Checksum	Link-state checksum (Fletcher checksum of the complete contents of the LSA).
Length	Length (in bytes) of the LSAs.
Router Priority	Interface priority of originating router.
Link Local Address	Link local address of the interface.
Number of Prefixes	Number of prefixes associated with the link.
Prefix Address and Length	List of prefixes associated with the link.
Options	LA indicates that the prefix is a local address. MC indicates that the prefix is multicast capable. NU indicates that the prefix is not unicast capable. P indicates that the prefix should be propagated at an NSSA area border.

The following is sample output from the **show ospfv3 database** command with the **network** keyword:

```
RP/0/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 78: show ospfv3 database network Field Descriptions

Field	Description
OSPFv3 Router with ID	Router ID number.
Process ID 1	OSPFv3 process name.
LS age	Link-state age.
Options	Type of service options (Type 0 only): DC—Supports demand circuits.E—Capable of processing external LSAs.MC—Forwards IP multicast.N—Supports Type 7 LSAs.R—Router is active.V6—Include in IPv6 routing calculations.
LS Type	Link-state type.
Link State ID	Link-state ID of the designated router.
Advertising Router	ID of the advertising router.
LS Seq Number	Link-state sequence (detects old or duplicate LSAs).
Checksum	Link-state checksum (Fletcher checksum of the complete contents of the LSA).
Length	Length (in bytes) of the LSA.
Attached Router	List of routers attached to the network, by router ID.

The following is sample output from the **show ospfv3 database** command with the **prefix** keyword:

```
RP/0/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 79: show ospfv3 database prefix Field Descriptions

Field	Description
OSPFv3 Router with ID	Router ID number.
Process ID 1	OSPFv3 process name.
LS age	Link-state age.
LS Type	Link-state type.
Link State ID	Link-state ID of the designated router.
Advertising Router	ID of the advertising router.
LS Seq Number	Link-state sequence (detects old or duplicate LSAs).
Checksum	Link-state checksum (Fletcher checksum of the complete contents of the LSA).
Length	Length (in bytes) of the LSA.
Referenced LSA Type	Router LSA or network LSA of the prefixes referenced.
Referenced Link State ID	Link-state ID of the router or network LSA.
Referenced Advertising Router	Advertising router of the referenced LSA.
Number of Prefixes	Number of prefixes listed in the LSA.
Prefix Address	Prefix associated with the router or network.
Prefix Length	Length of the prefix.
Options	LA indicates that the prefix is a local address. MC indicates that the prefix is multicast capable. NU indicates that the prefix is not unicast capable. P indicates the prefix should be propagated at an NSSA area border.

Field	Description
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 80: show ospfv3 database router Field Descriptions

Field	Description
OSPFv3 Router with ID	Router ID number.
Process ID 1	OSPFv3 process name.
LS age	Link-state age.
Options	Type of service options (Type 0 only): DC—Supports demand circuits.E—Capable of processing external LSAs.MC—Forwards IP multicast. N—Supports Type 7 LSAs.R—Router is active.V6—Include in IPv6 routing calculations.
LS Type	Link-state type.
Link State ID	Link-state ID of the designated router.

Field	Description
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.
Link connected to	The type of network to which this interface is connected. Values are: • Another Router (point-to-point). • A Transit Network. • A Virtual Link.
Link Metric	OSPF cost of using this link.
Local Interface ID	Number that uniquely identifies an interface on a router.

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*] [**vrf** {**all** *vrf-name*}]**flood-list** [*type interface-path-id*]

Syntax Description

process-name	(Optional) Name that uniquely identifies an OSPFv3 routing process. The process name is defined by the **router ospfv3** command. If this argument is included, only information for the specified routing process is displayed.
area-id	(Optional) Area number used to define the particular area.
type	Interface type. For more information, use the question mark (?) online help function.
interface-path-id	Physical interface or virtual interface. **Note** Use the **show interfaces** command to see a list of all interfaces currently configured on the router. For more information about the syntax for the router, use the question mark (?) online help function.
vrf	Specifies an OSPF VPN routing and forwarding (VRF) instance.
all	Display all the VRFs, excluding the default VRF.
vrf-name	Specifies the name of the of the OSPFv3 VRF.

Command Default

No default behavior or values

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

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 HundredGigE interface 0/3/0/0:

This table describes the significant fields shown in the display.

Table 81: 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.

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** [**vrf** {**all** *vrf-name*}][*type interface-path-id*]

Syntax Description

process-name	(Optional) Name that uniquely identifies an OSPFv3 routing process. The process name is defined by the **router ospfv3** command. If this argument is included, only information for the specified routing process is displayed.
area-id	(Optional) Area number used to define the particular area.
type	Interface type. For more information, use the question mark (?) online help function.
interface-path-id	Physical interface or virtual interface. **Note** Use the **show interfaces** command to see a list of all interfaces currently configured on the router. For more information about the syntax for the router, use the question mark (?) online help function.
vrf	Specifies an OSPF VPN routing and forwarding (VRF) instance.
all	Display all the VRFs, excluding the default VRF.
vrf-name	Specifies the name of the of the OSPFv3 VRF.

Command Default

No default behavior or values

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

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.

If you have configured BFD interface at the process level, then that applies to both area and interface level, unless area and interface level have different interval configurations specified.

Task ID

Task ID	Operations
ospf	read

Examples

The following is sample output from the **show ospfv3 interface** command when HundredGigE interface 0/2 5/0/0 is specified:

```
  RP/0/RP0/CPU0:router# show ospfv3 interface GigabitEthernet 0/
            2
            5

               /

            0
            0
            /0HundredGigE/25/0/0 is up, line protocol is up up
    Link Local address fe80::203213:a0ff1aff:fe9dfe14:f3fe3faa, Interface ID 2
    Area 0, Process ID 1foo, Instance ID 0, Router ID 101.01.01.2061
    Network Type BROADCAST, Cost: 101
    BFD enabled, interval 300 msec, multiplier 5
    Transmit Delay is 1 sec, State BDR, Priority 1
    Designated Router (ID) 102.02.02.2072, local address
fe80::20421b:c0ff53ff:fe22fe74:73feeab6
    Backup Designated router (ID) 101.01.01.2061, local address
fe80::203213:a0ff1aff:fe9dfe14:f3fe3faa
    Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
      Hello due in 00:00:06
    Index 0/21/1, flood queue length 0
    Next 0(0)/0(0)/0(0)
    Last flood scan length is 21, maximum is 92
    Last flood scan time is 0 msec, maximum is 1 0 msec
    Neighbor Count is 1, Adjacent neighbor count is 1
      Adjacent with neighbor 102.02.02.207 2  (Designated Router)
    Suppress hello for 0 neighbor(s)
```

This table describes the significant fields shown in the display.

Table 82: show ospfv3 interface Field Descriptions

Field	Description
HundredGigE	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.
BFD	BFD configuration state.
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.

Field	Description
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).
Neighbor Count	Count of network neighbors and list of adjacent neighbors.
Suppress hello	Count of neighbors suppressing hello messages.

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*] [**vrf** {**all** *vrf-name*}] **message-queue**

This command has no keywords or arguments.

Syntax Description

vrf	Specifies an OSPF VPN routing and forwarding (VRF) instance.
all	Display all the VRFs, excluding the default VRF.
vrf-name	Specifies the name of the of the OSPFv3 VRF.

Command Default

None

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

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*] [**vrf** {**all** *vrf-name*}]**neighbor** [*type interface-path-id*] [*neighbor-id*] [**detail**]

Syntax Description

process-name	(Optional) Name that uniquely identifies an OSPFv3 routing process. The process name is defined by the **router ospfv3** command. If this argument is included, only information for the specified routing process is displayed.
area-id	(Optional) Area ID. If you do not specify an area, all areas are displayed.
type	Interface type. For more information, use the question mark (?) online help function.
interface-path-id	Physical interface or virtual interface. **Note** Use the **show interfaces** command to see a list of all interfaces currently configured on the router. For more information about the syntax for the router, use the question mark (?) online help function.
neighbor-id	(Optional) Neighbor router ID.
detail	(Optional) Displays all neighbors given in detail (lists all neighbors).
vrf	Specifies an OSPF VPN routing and forwarding (VRF) instance.
all	Display all the VRFs, excluding the default VRF.
vrf-name	Specifies the name of the of the OSPFv3 VRF.

Command Default

No default behavior or values

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

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               HundredGigE 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 83: 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 HundredGigE 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 84: 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 HundredGigE 0/3/0/1 10.0.0.207

  Neighbors for OSPFv3 1

   Neighbor 10.0.0.207
```

```
    In the area 0 via interface HundredGigE 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 85: show ospfv3 neighbor HundredGigE 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 HundredGigE 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               HundredGigE 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 86: show ospfv3 neighbor HundredGigE 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 HundredGigE interface 0/3/0/1:

```
RP/0/RP0/CPU0:router# show ospfv3 neighbor HundredGigE 0/3/0/1 detail

  Neighbors for OSPFv3 1

   Neighbor 10.0.0.207
      In the area 0 via interface HundredGigE 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 87: show ospfv3 neighbor HundredGigE 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.

The following is sample output showing neighbor information with BFD enabled for HundredGigE interface 0/5/0/0:

```
RP/0/RP0/CPU0:router# show ospfv3 neighbor detail
Thu Sep 11 02:02:46.267 UTC

  Neighbors for OSPFv3 foo

   Neighbor 2.2.2.2
      In the area 0 via interface HundredGigE 0/5/0/0     BFD enabled
      Neighbor: interface-id 2, link-local address fe80::21b:53ff:fe74:eab6
      Neighbor priority is 1, State is FULL, 6 state changes
      DR is 2.2.2.2 BDR is 1.1.1.1
      Options is 0x13
      Dead timer due in 00:00:32
      Neighbor is up for 00:06:16
      Index 1/1/1, 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: 1
```

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*] [**vrf** {**all** *vrf-name*}]**request-list** [*type interface-path-id*] [*neighbor-id*]

Syntax Description

process-name	(Optional) Name that uniquely identifies an OSPFv3 routing process. The process name is defined by the **router ospfv3** command. If this argument is included, only information for the specified routing process is displayed.
area-id	(Optional) Area ID. If you do not specify an area, all areas are displayed.
type	(Optional) Interface type. For more information, use the question mark (?) online help function.
interface-path-id	(Optional) Physical interface or virtual interface. **Note** Use the **show interfaces** command to see a list of all interfaces currently configured on the router. For more information about the syntax for the router, use the question mark (?) online help function.
neighbor-id	(Optional) Router ID of the OSPFv3 neighbor. This argument must be in 32-bit dotted-decimal notation, similar to an IPv4 address.
vrf	Specifies an OSPF VPN routing and forwarding (VRF) instance.
all	Display all the VRFs, excluding the default VRF.
vrf-name	Specifies the name of the of the OSPFv3 VRF.

Command Default

No default behavior or values

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

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 HundredGigE 0/3/0/0

  Request Lists for OSPFv3 1

   Neighbor 10.0.0.207, interface HundredGigE 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 88: 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.

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*] [**vrf** {**all** *vrf-name*}]**retransmission-list** [*type interface-path-id*] [*neighbor-id*]

Syntax Description

process-name	(Optional) Name that uniquely identifies an Open Shortest Path First Version 3 (OSPFv3) routing process. The process name is defined by the **router ospfv3** command. If this argument is included, only information for the specified routing process is displayed.
area-id	(Optional) Area ID. If you do not specify an area, all areas are displayed.
type	(Optional) Interface type. For more information, use the question mark (?) online help function.
interface-path-id	(Optional) Physical interface or virtual interface. **Note** Use the **show interfaces** command to see a list of all interfaces currently configured on the router. For more information about the syntax for the router, use the question mark (?) online help function.
neighbor-id	(Optional) IP address of the OSPFv3 neighbor.
vrf	Specifies an OSPF VPN routing and forwarding (VRF) instance.
all	Display all the VRFs, excluding the default VRF.
vrf-name	Specifies the name of the of the OSPFv3 VRF.

Command Default

No default behavior or values

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

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 HundredGigE interface 0/3/0/0:

```
RP/0/RP0/CPU0:router#show ospfv3 retransmission-list 10.0.124.4 HundredGigE 0/3/0/0

  Neighbor 10.0.124.4, interface HundredGigE 0/3/0/0 address fe80::3034:30ff:fe33:3742
```

This table describes the significant fields shown in the display.

Table 89: show ospfv3 retransmission-list 10.0.124.4 HundredGigE 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.

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*] [**vrf** {**all** *vrf-name*}]**routes** [**external** | **connected**] [*ipv6-prefix*/*prefix-length*]
show ospfv3 [*process-name*] [**vrf** {**all** *vrf-name*}]**routes summary**

Syntax Description

process-name	(Optional) Name that uniquely identifies an OSPFv3 routing process. The process name is defined by the **router ospf** command. If this argument is included, only information for the specified routing process is displayed.
external	(Optional) Displays routes redistributed from other protocols.
connected	(Optional) Displays connected routes.
ipv6-prefix	(Optional) IP Version 6 (IPv6) prefix, which limits output to a specific route. This argument must be in the form documented in RFC 2373, in which the address is specified in hexadecimal using 16-bit values between colons.
/ *prefix-length*	(Optional) Length of the IPv6 prefix. A decimal value that indicates how many of the high-order contiguous bits of the address compose the prefix (the network portion of the address). A slash must precede the decimal value.
summary	Displays a summary of the route table.
vrf	Specifies an OSPF VPN routing and forwarding (VRF) instance.
all	Display all the VRFs, excluding the default VRF.
vrf-name	Specifies the name of the of the OSPFv3 VRF.

Command Default

No default behavior or values

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

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
      HundredGigE 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
  HundredGigE 0/3/0/0, fe80::3034:30ff:fe33:3742
      10.0.0.207/1
  * 3333::/56, Ext2, cost 20/1, P:0 F:0
  HundredGigE 0/3/0/0, fe80::3034:30ff:fe33:3742
      10.0.0.207/0
  * 6050::/56, Ext2, cost 20/1, P:0 F:0
  HundredGigE 0/3/0/0, fe80::3034:30ff:fe33:3742
      10.0.0.207/1
  * 7002::/56, Intra, cost 10/0, area 0
      HundredGigE 0/0/0/0, connected

  * 3000:11:22::/64, Inter, cost 21/0, area 1
  HundredGigE 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 90: 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.
HundredGigE 0/3/0/0	Packets destined for prefix 3000:11:22::/64 are sent over the HundredGigE 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.

show ospfv3 statistics rib-thread

To display RIB thread statistics, use the **show ospfv3 statistics rib-thread** command in XR EXEC mode.

show ospfv3 [*process-name* [*area-id*]] **statistics rib-thread**

Syntax Description

process-name	(Optional) Name that uniquely identifies an OSPF routing process. The process name is defined by the **router ospfv3** command. If this argument is included, only information for the specified routing process is displayed.
area id	(Optional) Area number used to define the particular area.

Command Default

None

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

Task ID

Task ID	Operation
ospf	read

This is sample output from **show ospfv3 statistics rib-thread** command:

```
RP/0/RP0/CPU0:router#show ospfv3 0 statistics rib-thread
Mon May 10 17:48:29.011 CEST
OSPFv3 0 RIB thread statistics
 Queue statistics:
  Last entry dequeue           10127056 msecs ago (14:59:42.171)
  RIB thread active            NO
  Total RIB thread signals     30
  Current queue length         0
  Maximum queue length         2
  Total entries queued         31
  Total entries dequeued       31
  Maximum latency (msec)       5.000
  Average latency (msec)       0.323
 Queue errors:
  Enqueue errors               0
  Dequeue errors               0
 RIB batch statistics:
  Batches sent to RIB          31
  Batch all routes OK          31
  Batch some routes backup     0
 RIB batch errors:
  Batches version mismatch     0
  Batches missing connection   0
```

```
Batches no table             0
Batch route table limit      0
Batch route errors           0
Batch errors                 0
Route table limit            0
Route path errors            0
Route errors                 0
Path table limit             0
Path errors                  0
```

show ospfv3 summary-prefix

To display Open Shortest Path First Version 3 (OSPFv3) aggregated summary address information, use the **show ospfv3 summary-prefix** command in XR EXEC mode.

show ospfv3 [*process-name*] [**vrf** {**all** *vrf-name*}]**summary-prefix**

Syntax Description

process-name	(Optional) Name that uniquely identifies an OSPFv3 routing process. The process name is defined by the **router ospfv3** command. If this argument is included, only information for the specified routing process is displayed.
vrf	Specifies an OSPF VPN routing and forwarding (VRF) instance.
vrf-name	Specifies the name of the of the OSPFv3 VRF.

Command Default

No default behavior or values

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

Use the **show 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 91: 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.

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*] [**vrf** {**all** *vrf-name*}]**virtual-links**

Syntax Description

process-name	(Optional) Name that uniquely identifies an OSPFv3 routing process. The process name is defined by the **router ospfv3** command. If this argument is included, only information for the specified routing process is displayed.
vrf	Specifies an OSPF VPN routing and forwarding (VRF) instance.
vrf-name	Specifies the name of the of the OSPFv3 VRF.

Command Default

No default behavior or values

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

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 HundredGigE 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 92: 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.

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**] [**all***protocol*]

Syntax Description

afi-all	(Optional) Specifies all address families.
ipv4	(Optional) Specifies an IPv4 address family.
ipv6	(Optional) Specifies an IPv6 address family.
all	(Optional) Specifies all protocols for a given address family.
protocol	(Optional) Specifies a routing protocol. For the IPv4 address family, the options are: • **bgp** • **isis** • **ospf** • **rip** For the IPv6 address family, the options are: • **bgp** • **isis** • **ospfv3**

Command Default

The default address family is IPv4.

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

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
  HundredGigE 0/2/0/2
  Loopback1
```

This table describes the significant fields shown in the display.

Table 93: show protocols Field Descriptions

Field	Description
Router Id	Router ID of the OSPFv3 process.
Distance	Administrative distance for the protocol. This distance determines the priority the Routing Information Base (RIB) gives to the routes, as opposed to other protocols, for example, IS-IS.
Redistribution	Protocols from which this OSPFv3 process is redistributing routes.
Area	OSPFv3 areas defined in this process, followed by their associated interfaces.

snmp context (OSPFv3)

To specify an SNMP context for an OSPFv3 instance, use the **snmp context** command in router configuration mode or in VRF configuration mode. To remove the SNMP context, use the **no** form of this command.

snmp context *context_name*

Syntax Description

context_name	Specifies name of the SNMP context for OSPFv3 instance.

Command Default

SNMP context is not specified.

Command Modes

Router OSPFv3 configuration

VRF configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

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

The snmp-server commands need to be configured to perform SNMP request for the OSPF instance. Refer *SNMP Server Commands* module in *System Management Command Reference* for information on using the snmp-server commands.

Note To map an SNMP context with a protocol instance, topology or VRF entity, use the **snmp-server context mapping** command. However, the **feature** option of this command does not work with OSPFv3 protocol.

Task ID

Task ID	Operation
ospf	read, write

This example shows how to configure an SNMP context *foo* for OSPFv3 instance *100*:

```
RP/0/RP0/CPU0:router#configure
RP/0/RP0/CPU0:router(config)#router ospfv3 100
RP/0/RP0/CPU0:router(config-ospf)#snmp context foo
```

This example shows how to configure **snmp-server** commands to be used with the **snmp context** command:

```
RP/0/RP0/CPU0:router(config)#snmp-server host 10.0.0.2 traps version 2c public udp-port
```

```
1620
RP/0/RP0/CPU0:router(config)#snmp-server community public RW
RP/0/RP0/CPU0:router(config)#snmp-server contact foo
RP/0/RP0/CPU0:router(config)#snmp-server community-map public context foo
```

This is a sample SNMP context configuration for OSPFv3 instance *100*:

```
snmp-server host 10.0.0.2 traps version 2c public udp-port 1620
snmp-server community public RW
snmp-server contact foo

snmp-server community-map public context foo


router ospfv3 100
 router-id 2.2.2.2
 bfd fast-detect
 nsf cisco
snmp context foo
 area 0
  interface Loopback1
  !
 !
 area 1
  interface  0/2/0/1
   demand-circuit enable
  !
  interface POS0/3/0/0
  !
  interface POS0/3/0/1
  !
 !
!
```

snmp trap (OSPFv3)

To enable SNMP trap for an OSPFv3 instance, use the **snmp trap** command in VRF configuration mode. To disable SNMP trap for the OSPFv3 instance, use the **no** form of this command.

snmp trap

Syntax Description

This command has no keywords or arguments.

Command Default

Disabled.

Command Modes

OSPFv3 VRF configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

Task ID

Task ID	Operation
ospf	read, write

This example shows how to enable SNMP trap for OSPFv3 instance *100* under VRF *vrf-1*:

```
RP/0/RP0/CPU0:router#configure
RP/0/RP0/CPU0:router(config)#router ospfv3 100
RP/0/RP0/CPU0:router(config-ospf)#vrf vrf-1
RP/0/RP0/CPU0:router(config-ospf-vrf)#snmp trap
```

snmp trap rate-limit (OSPFv3)

To control the number of traps that OSPFv3 sends by configuring window size and the maximum number of traps during that window, use the **snmp trap rate-limit** command in router OSPFv3 configuration mode or OSPFv3 VRF configuration mode. To disable configuring the window size and maximum number of traps during the window, use the **no** form of this command.

snmp trap rate-limit *window-size max-num-traps*

Syntax Description

window-size	Specifies the trap rate limit sliding window size. The range is 2 to 60 windows.
max-num-traps	Specifies the maximum number of traps sent in window time. The range is 0 to 300 traps.

Command Default

None

Command Modes

Router OSPFv3 configuration

OSPFv3 VRF configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

Task ID

Task ID	Operation
ospf	read, write

This example shows how to set the trap rate limit sliding window size to 50 and the maximum number of traps sent to 250 for OSPFv3 instance *100* under vrf *vrf1*:

```
RP/0/RP0/CPU0:router#configure
RP/0/RP0/CPU0:router(config)#router ospfv3 100
RP/0/RP0/CPU0:router(config-ospfv3)#vrf vrf1
RP/0/RP0/CPU0:router(config-ospfv3-vrf)#snmp trap rate-limit 50 250
```

spf prefix-priority (OSPFv3)

To prioritize OSPFv3 prefix installation into the global Routing Information Base (RIB) during Shortest Path First (SPF) run, use the **spf prefix-priority** command in XR Config mode or VRF configuration mode. To return to the system default value, use the **no** form of this command.

spf prefix-priority route-policy *policy-name* [**disable**]
spf prefix-priority route-policy *policy-name*

Syntax Description

route-policy	Specifies the route-policy to prioritize route installation.
policy-name	Name of the route policy.
disable	Disables SPF prefix priority

Command Default

SPF prefix prioritization is disabled.

Command Modes

XR Config mode

VRF configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

Task ID

Task ID	Operation
ospf	read, write

This example shows how to configure OSPFv3 SPF prefix prioritization:

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# prefix-set ospf3-critical-prefixes
RP/0/RP0/CPU0:router(config-pfx)# 66.0.0.0/16
RP/0/RP0/CPU0:router(config-pfx)# end-set
RP/0/RP0/CPU0:router(config)# route-policy ospf3-spf-priority
RP/0/RP0/CPU0:router(config-rpl)# if destination in ospf-critical-prefixes then set
spf-priority critical
endif
RP/0/RP0/CPU0:router(config-rpl)# end-policy
RP/0/RP0/CPU0:router(config-rpl)# commit
RP/0/RP0/CPU0:router(config-rpl)# exit
RP/0/RP0/CPU0:router(config)# router ospfv3 1
RP/0/RP0/CPU0:router(config-ospf)# router-id 66.0.0.1
RP/0/RP0/CPU0:router(config-ospf)# spf prefix-priority route-policy ospf-spf-priority
```

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**]

Syntax Description

no-summary	(Optional) Prevents an area border router (ABR) from sending summary link advertisements into the stub area. Areas with this option are known as *totally stubby* areas.

Command Default

No stub area is defined.

Command Modes

Area configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

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
```

stub-router

To modify self originated router LSAs when stub router is active, use the **stub-router** command in an appropriate configuration mode. To disable this function, use the **no** form of this command.

stub-router router-lsa [**r-bit** | **v6-bit** | **max-metric**] [**always**] [**on-proc-migration** *interval*] [**on-proc-restart** *interval*] [**on-switchover** *interval*] [**on-startup** [*interval* | **wait-for-bgp**]] [**summary-lsa** [*metric*]] [**external-lsa** [*metric*]] [**include-stub**]
stub-router router-lsa [**r-bit** | **v6-bit** | **max-metric**]

Syntax Description

router-lsa	Specifies that always originate router link-state advertisements (LSAs) with the stub-router.
r-bit	Router-LSAs are originated with R-bit clear (v6 bit set), which means the node does not act as a transit router. Directly connected networks (native to OSPF) are still reachable within the OSPF area.
v6-bit	Router-LSAs are originated with V6 bit clear (and also r-bit clear). That means the node is not willing to receive any ipv6 traffic. Other ospfv3 routers won't install any route to a node with v6-bit clear.
max-metric	Router-LSAs are originated with maximum metric. Unlike the r-bit and v6-bit mode, the router may still act as a transit node, if there is no alternate path.
always	Stub-router mode is activated unconditionally.
on-proc-migration	Stub-router mode is activated for the desired period of time, upon ospfv3 process migration.
on-proc-restart	Stub-router mode is activated for the desired period of time, upon ospfv3 process restart.
on-switchover	Stub-router mode is activated for the desired period of time, upon RP failover.
on-startup	Stub-router mode is activated (for configured time, or until BGP converges) upon router startup (boot).
wait-for-bgp	Stub-router mode is terminated upon BGP convergence in ipv6 unicast address family. This option could only be used in the global routing table, not in a non-default VRF. This option is only supported with the on-startup trigger when the router boots.
summary-lsa	If enabled, summary LSAs are advertised with modified metric when stub-router is active. This configuration is applicable to max-metric mode. In r-bit mode, ABR/ASBR functionality is implicitly disabled and routers will not use this node as an ABR/ASBR, since it declares no transit capability (r-bit clear). If enabled and metric is not explicitly configured, the default metric for summary LSAs when stub-router active is 16711680 (0xFF0000).

external-lsa	If enabled, external LSAs are advertised with modified metric when stub-router is active. This configuration is applicable to max-metric mode. In r-bit mode, ABR/ASBR functionality is implicitly disabled and routers will not use this node as an ABR/ASBR, since it declares no transit capability (r-bit clear). If enabled and metric is not explicitly configured, the default metric for external LSAs when stub-router active is 16711680 (0xFF0000).
include-stub	If enabled, intra-area-prefix LSAs that are referencing router LSA are advertised with maximum metric (0xffff) when stub-router is active. Intra-area-prefix LSAs that are referencing network LSA do not change metric Can be used in r-bit and max-metric modes. /128 prefixes that are normally advertised with LA-bit set and 0 metric are also advertised with maximum metric and LA-bit clear when stub-router is active.

Command Default

Disabled.

Command Modes

Router OSPFv3 configuration

OSPFv3 VRF configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

Only one method (r-bit, v6-bit, max-metric) could be activated at a time. Configuring the methods simultaneously, or different method per trigger, is not supported.

Task ID

Task ID	Operation
ospf	read, write

This example shows how to configure router LSAs are originated with R-bit clear under OSPFv3 VRF, *vrf_1*:

```
RP/0/RP0/CPU0:router#configure
RP/0/RP0/CPU0:router(config)#router ospfv3 1
RP/0/RP0/CPU0:router(config-ospfv3)#vrf vrf_1
RP/0/RP0/CPU0:router(config-ospfv3-vrf)#stub-router router-lsa r-bit
```

summary-prefix (OSPFv3)

To create aggregate addresses for routes being redistributed from another routing protocol into Open Shortest Path First Version 3 (OSPFv3) protocol, use the **summary-prefix** command in an appropriate configuration mode. To stop summarizing redistributed routes, use the **no** form of the command.

summary-prefix *ipv6-prefix/prefix-length* [**not-advertise**] **tag** *tag*

Syntax Description

ipv6-prefix	Summary prefix designated for a range of IP Version 6 (IPv6) prefixes. This argument must be in the form documented in RFC 2373, where the address is specified in hexadecimal using 16-bit values between colons.
/ *prefix-length*	Length of the IPv6 prefix. A decimal value that indicates how many of the high-order contiguous bits of the address compose the prefix (the network portion of the address). A slash must precede the decimal value.
not-advertise	(Optional) Suppresses summary routes that match the address and mask pair from being advertised.
tag *tag*	(Optional) Specifies a tag value that can be used as a "match" value for controlling redistribution.

Command Default

When this command is not used in router configuration mode, aggregate addresses are not created for routes being redistributed from another routing protocol into the OSFPv3 protocol.

Command Modes

Router OSPFv3 configuration

OSPFv3 VRF configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

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
```

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*

Syntax Description

milliseconds	Minimum delay (in milliseconds) that must pass between acceptance of the same LSA arriving from neighbors. Range is 0 to 60000 milliseconds.

Command Default

1000 milliseconds

Command Modes

Router OSPFv3 configuration

OSPFv3 VRF configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

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
```

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*

Syntax Description

milliseconds	Time (in milliseconds) at which LSAs in the flooding queue are paced in between updates. Range is 5 milliseconds to 100 milliseconds.

Command Default

milliseconds: 33

Command Modes

Router OSPFv3 configuration

OSPFv3 VRF configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

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
```

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*

Syntax Description

seconds	Interval (in seconds) at which LSAs are grouped and refreshed, checksummed, or aged. Range is 10 to 1800 seconds.

Command Default

seconds : 240

OSPFv3 LSA group pacing is enabled by default.

Command Modes

Router OSPFv3 configuration

OSPFv3 VRF configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

Use the **timers pacing lsa-group** command to control the rate at which LSA updates occur so that high CPU or buffer utilization that can occur when an area is flooded with a very large number of LSAs can be reduced. The default settings for OSPFv3 packet pacing timers are suitable for the majority of deployments. Do not change the packet pacing timers unless all other options to meet OSPFv3 packet flooding requirements have been exhausted. Specifically, network operators should prefer summarization, stub area usage, queue tuning, and buffer tuning before changing the default flooding timers. Furthermore, no guidelines exist for changing timer values; each OSPFv3 deployment is unique and should be considered on a case-by-case basis. The network operator assumes the risks associated with changing the default timer values.

Cisco IOS XR software groups the periodic refresh of LSAs to improve the LSA packing density for the refreshes in large topologies. The group timer controls the interval used for group refreshment of LSAs; however, this timer does not change the frequency that individual LSAs are refreshed (the default refresh rate is every 30 minutes).

The duration of the LSA group pacing is inversely proportional to the number of LSAs the router is handling. For example, if you have about 10,000 LSAs, decreasing the pacing interval would benefit you. If you have a very small database (40 to 100 LSAs), increasing the pacing interval to 10 to 20 minutes might benefit you slightly.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to configure OSPFv3 group packet-pacing updates between LSA groups to occur in 60-second intervals for OSPFv3 routing process 1:

```
RP/0/RP0/CPU0:router(config)# router ospfv3 1
RP/0/RP0/CPU0:router(config-ospfv3)# timers pacing lsa-group 60
```

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*

Syntax Description

milliseconds	Time (in milliseconds) at which LSAs in the retransmission queue are paced. Range is 5 milliseconds to 100 milliseconds.

Command Default

milliseconds : 66

Command Modes

Router OSPFv3 configuration

OSPFv3 VRF configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

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
```

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*

Syntax Description

start-interval	Minimum delay (in milliseconds) for the generation of LSAs. The first instance of LSA is always generated immediately upon a local OSPFv3 topology change. The generation of the next LSA is not before the start interval. Range is 0 to 600000 milliseconds.
hold-interval	Incremental time (in milliseconds). This value is used to calculate the subsequent rate limiting times for LSA generation. Range is 1 to 600000 milliseconds.
max-interval	Maximum wait time (in milliseconds) between generation of the same LSA. Range is 1 to 600000 milliseconds.

Command Default

start-interval : 500 milliseconds 50 milliseconds

hold-interval : 5000 milliseconds 200 milliseconds

max-interval : 5000 milliseconds

Command Modes

Router OSPFv3 configuration

OSPFv3 VRF configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

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
```

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*

Syntax Description

spf-start	Initial SPF schedule delay (in milliseconds). Range is 1 to 600000 milliseconds.
spf-hold	Minimum hold time (in milliseconds) between two consecutive SPF calculations. Range is 1 to 600000 milliseconds.
spf-max-wait	Maximum wait time (in milliseconds) between two consecutive SPF calculations. Range is 1 to 600000 milliseconds.

Command Default

spf-start : 5000 milliseconds50 milliseconds

spf-hold: 10000 milliseconds200 milliseconds

spf-max-wait: 10000 milliseconds5000 milliseconds

Command Modes

Router OSPFv3 configuration

OSPFv3 VRF configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

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*

Syntax Description

size	Deletes existing buffer and creates one with *N* entries.
buffer_name	Specifies a buffer from one of the 15 listed buffers. Refer Table 1 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 7.0.12	This command was introduced.

Usage Guidelines

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 94: Buffer Types

Name	Description
adj	adjacency
adj_cycle	dbd/flood events/pkts
config	config events
errors	errors
events	mda/rtrid/bfd/vrf
ha	startup/HA/NSF
hello	hello events/pkts
idb	interface
pkt	I/O packets
rib	rib batching
spf	spf/topology

Name	Description
spf_cycle	spf/topology detail
te	mpls-te
test	testing info
mq	message queue info

Task ID

Task ID	Operations
ospf	read, write

Examples

This example shows how to set 1024 error trace entries:

```
RP/0/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*

Syntax Description

seconds	Time (in seconds) required to send a link-state update. Range is 1 to 65535 seconds.

Command Default

1 second

Command Modes

Process configuration

Area configuration

Interface configuration

Virtual-link configuration

OSPFv3 VRF configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

Link-state advertisements (LSAs) in the update packet must have their ages incremented by the amount specified in the *seconds* argument before transmission. The value assigned should take into account the transmission and propagation delays for the interface.

If the delay is not added before transmission over a link, the time in which the LSA propagates over the link is not considered. This setting has more significance on very low-speed links.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to configure a transmit delay of 3 seconds for HundredGigE interface 0/3/0/0:

```
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 HundredGigE 0/3/0/0
RP/0/RP0/CPU0:router(config-ospfv3-ar-if)# transmit-delay 3
```

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*

Syntax Description

router-id	Router ID associated with the virtual link neighbor. The router ID appears in the **show ospfv3** display. This value must be entered in 32-bit dotted-decimal notation, similar to an IP Version 4 (IPv4) address. There is no default.

Command Default

No virtual links are defined.

Command Modes

Area configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

In OSPFv3, when there exists a path through another non-backbone area over which the virtual link can function, all areas must be connected to a backbone area. If the connection to the backbone is lost, it can be repaired by establishing a virtual link.

Virtual links, which are defined in the submode of the area they transit, are in effect virtual point-to-point interfaces belonging to area 0 (the backbone). The virtual links inherit parameter values from the backbone area, rather than the transit area in which they are defined.

Each virtual link neighbor must include the router ID of the virtual link neighbor for the link to be properly established. Use the **show ospfv3** command to display the router ID of an OSPFv3 process.

Use the **virtual-link** command to place the router in virtual-link configuration mode (config-router-ar-vl), from which you can configure virtual-link-specific settings. Commands configured under this mode (such as the **transmit-delay** command) are automatically bound to that virtual link.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to establish a virtual link with default values for all optional parameters:

```
RP/0/RP0/CPU0:router(config)# router ospfv3 201
RP/0/RP0/CPU0:router(config-ospfv3)# area 1
RP/0/RP0/CPU0:router(config-ospfv3-ar)# virtual-link 10.3.4.5
```

vrf (OSPFv3)

To configure an Open Shortest Path First Version 3 (OSPFv3) VPN routing and forwarding (VRF) instance and enter OSPFv3 VRF configuration submode, use the **vrf** command in XR Config mode. To terminate an OSPFv3 VRF, use the **no** form of this command.

vrf *vrf-name*

Syntax Description

vrf-name	Specifies a name for the OSPFV3 vrf. If a name is not specified, the default vrf is assumed.

Command Default

No OSPFv3 VRF is configured.

Command Modes

XR Config mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

Use the **vrf** command to explicitly configure a VRF. This command creates a separate instance of OSPFv3 within the process. Commands configured under the VRF configuration mode (such as the **interface** [OSPFv3] and **authentication** [OSPFv3] commands) are automatically bound to that VRF.

To modify or remove the VRF, the *vrf-name* argument format must be the same as the format used when creating the VRF.

Task ID

Task ID	Operation
ospf	read, write

This example shows how to configure VRF *vrf_1* and enter OSPFv3 VRF configuration submode:

```
RP/0/RP0/CPU0:router#configure
RP/0/RP0/CPU0:router(config)#router ospfv3 osp3
RP/0/RP0/CPU0:router(config-ospfv3)#vrf vrf_1
RP/0/RP0/CPU0:router(config-ospfv3-vrf)#
```



RIB Commands

- address-family next-hop dampening disable, on page 573
- clear route, on page 574
- maximum prefix (RIB), on page 576
- lcc, on page 577
- rcc, on page 578
- recursion-depth-max, on page 579
- router rib, on page 580
- rump always-replicate, on page 581
- show lcc statistics, on page 582
- show rcc, on page 584
- show rcc statistics, on page 586
- show rcc vrf, on page 588
- show rib, on page 589
- show rib afi-all, on page 591
- show rib attributes, on page 593
- show rib client-id, on page 594
- show rib clients, on page 596
- show rib extcomms, on page 598
- show rib firsthop, on page 600
- show rib history, on page 602
- show rib next-hop, on page 604
- show rib opaques, on page 607
- show rib protocols, on page 609
- show rib recursion-depth-max, on page 611
- show rib statistics, on page 613
- show rib tables, on page 615
- show rib trace, on page 617
- show rib vpn-attributes, on page 619
- show rib vrf, on page 621
- show route, on page 622
- show route backup, on page 629
- show route best-local, on page 632
- show route connected, on page 634

- show route local, on page 636
- show route longer-prefixes, on page 638
- show route next-hop, on page 640
- show route quarantined, on page 642
- show route resolving-next-hop, on page 644
- show route static, on page 646
- show route summary, on page 648
- table ip-only activate vrf, on page 651
- tag-map tag, on page 652

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 XR Config mode. To enable RIB next-hop dampening, use the **no** form of this command.

address-family { **ipv4** | **ipv6** } **next-hop dampening disable**

Syntax Description

ipv4	Specifies IP Version 4 (IPv4) address prefixes.
ipv6	Specifies IP Version 6 (IPv6) address prefixes.

Command Default

RIB next-hop dampening is enabled.

Command Modes

XR Config mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

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 [**vrf** {*vrf-name* | **all**}] {**ipv4** | **ipv6** | **afi-all** | **safi-all**} {**unicast** | **multicast** | **safi-all**} [**topology** *topo-name*] [*ip-address mask*]

Syntax Description

vrf { *vrf-name* | **all** }	(Optional) Specifies a particular VPN routing and forwarding (VRF) instance or all VRF instances.
ipv4	Specifies IP Version 4 address prefixes.
ipv6	Specifies IP Version 6 address prefixes.
afi -all **safi -all**	Specifies IP Version 4 and IP Version 6 address prefixes.
safiunicast	Specifies unicast address prefixes.
multicast	Specifies multicast address prefixes.
safi-all	Specifies unicast and multicast address prefixes.
topology *topo-name*	(Optional) Specifies topology table information and name of the topology table.
ip-address node-id	(Optional) Clears hardware resource counters from the designated node. The *node-id* argument is entered in the *rack/slot/module* notation.
ip-address	Network IP address about which routing information should be displayed.
mask	Network mask specified in either of two ways: Network mask can be a four-part, dotted-decimal address. For example, 255.0.0.0 indicates that each bit equal to 1 means the corresponding address bit is a network address. Network mask can be indicated as a slash (/) and number. For example, /8 indicates that the first 8 bits of the mask are 1s, and the corresponding bits of the address are the network address.

Command Default

If a **vrf** *vrf-name* is not specified, routes are cleared from the default IPv4 unicast VRF.

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

Use the **clear route** command to clear routes from an IP routing table to a specific network, a matching subnet address, or all routes.

The topology keyword must be accompanied by the ipv4 multicast keywords, except when the afi-all keyword or the safi-all keyword is specified.

Task ID

Task ID	Operations
rib	read, 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
```

maximum prefix (RIB)

To set the prefix limit for the VPN routing and forwarding (VRF) instance, use the **maximum prefix** command in global VRF address family configuration mode. To set the prefix limits to the default values, use the **no** form of this command.

maximum prefix *maximum* [*mid-threshold*]

Syntax Description

maximum	Maximum number of prefixes allowed in the VRF instance. Range is 32 to 2000000.
mid-threshold	(Optional) Integer specifying at what percentage of the *maximum* argument value the software starts to generate a Simple Network Management Protocol (SNMP) trap. Range is 1 to 100.

Command Default

No default behavior or values

Command Modes

Global VRF address family configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

Use the **maximum prefix** command to configure a maximum number of prefixes that a VRF instance is allowed to receive.

Task ID

Task ID	Operations
rib	read, write

Examples

The following example shows how to set the maximum number of prefixes allowed to 1000:

```
RP/0/RP0/CPU0:router(config)# vrf vrf-A
RP/0/RP0/CPU0:router(config-vrf)# address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-vrf-af)#  maximum prefix 1000
```

A maximum number of routes is applicable to dynamic routing protocols as well as static or connected routes. When maximum prefix is configured, an **syslog** message is generated in the following conditions:

1. if the number of routes is above the threshold when“maximum prefix” configuration has been committed
2. if the number routes reaches the configured “maximum prefix” values for the VRF

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

ipv4	Specifies IP Version 4 address prefixes.
ipv6	Specifies IP Version 6 address prefixes.
unicast	Specifies unicast address prefixes.
period *milliseconds*	Specifies the period between scans in milliseconds. Range is 100 to 600000 milliseconds.

Command Default

Label consistency checker is disabled.

Command Modes

XR Config mode

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

Task ID

Task ID	Operation
ipv4	read, write
ipv6	read, write

This example shows how to enbale lcc for IPv6 lables:

```
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* }

Syntax Description

ipv4	Specifies IP Version 4 address prefixes.
ipv6	Specifies IP Version 6 address prefixes.
unicast	Specifies unicast address prefixes.
period *milliseconds*	Specifies the period between scans in milliseconds. Range is 100 to 600000 milliseconds.

Command Default

Route consistency checker is disabled.

Command Modes

XR Config mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

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 XR Config mode. To set the recursion checks to the default value, use the **no** form of this command.

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

XR Config mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

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

Syntax Description

This command has no arguments or keywords.

Command Default

Router configuration mode is not enabled.

Command Modes

XR Config mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

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 XR Config mode. To diable replication from uRIB to muRIB, use the **no** form of this command.

rump always-replicate [*access-list*]

Syntax Description

access-list-name	(Optional) Name of the access list.

Command Default

Replication from uRIB to muRIB is enabled.

Command Modes

XR Config mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

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 *scan-id-value*	Specifies the scan ID value. The range is between <0-100000>.
summary	Displays a summary of the BG route consistency check statistics.

Command Default

None

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

Task ID

Task ID	Operation
ipv4	read
ipv6	read

This example shows background scan statistics for AFI-SAFI mplsv6-unicast:

```
RP/0/RP0/CPU0:router#show lcc ipv6 unicast statistics

Background Scan Statistics for AFI-SAFI mplsv6-unicast:
=======================================================

Scan enabled:           False
Current scan-id:        0               Scan triggered:        False
Configured period:      60              Current period:        0

Paused by range scan: False
Paused by route churn: False
Paused by error scan: False

Last data sent: 0 entries               Damping percent:       70
Default route churn:    10              Current route churn:   0
Route churn last calculated at          Dec 31 16:00:00.000
```

```
Logs stored for background scan ids:

Log for AFI-SAFI mplsv6-unicast:
================================

End Of Logs
```

This example shows background scan statistics for AFI-SAFI mplsv4-unicast:

```
RP/0/RP0/CPU0:router#show lcc ipv4 unicast statistics


Background Scan Statistics for AFI-SAFI mplsv4-unicast:
=======================================================

Scan enabled:           False
Current scan-id:        0               Scan triggered:        False
Configured period:      60              Current period:        0

Paused by range scan: False
Paused by route churn: False
Paused by error scan: False

Last data sent: 0 entries               Damping percent:       70
Default route churn:    10              Current route churn:   0
Route churn last calculated at          Dec 31 16:00:00.000

Logs stored for background scan ids:

Log for AFI-SAFI mplsv4-unicast:
================================

End Of Logs
```

show rcc

To display route consistency checker (RCC) information, use the **show rcc** command in XR EXEC mode.

show rcc {**ipv4** | **ipv6**} **unicast** [*prefix netmask* **vrf** *vrf-name*]

Syntax Description

ipv4	Specifies IP Version 4 address prefixes.
ipv6	Specifies IP Version 6 address prefixes.
unicast	Specifies unicast address prefixes.
prefix	(Optional) Starting prefix.
netmask	(Optional) Network mask.
vrf *vrf-name*	(Optional) Specifies a particular VPN routing and forwarding (VRF) instance or all VRF instances.

Command Default

No default behavior or values

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

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 *scan-id-value*	Specifies the scan ID value. The range is between <0-100000>.
summary	Displays a summary of the BG route consistency check statistics.

Command Default

None

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

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               Scan triggered:        False
Configured period:      60              Current period:        0

Paused by range scan: False
Paused by route churn: False
Paused by error scan: False

Last data sent: 0 entries               Damping percent:       70
Default route churn:    10              Current route churn:   0
Route churn last calculated at          Dec 31 16:00:00.000
```

```
Logs stored for background scan ids:

Log for AFI-SAFI ipv6-unicast:
================================

End Of Logs
```

This example shows background scan statistics for AFI-SAFI Ipv4 unicast:

```
RP/0/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 rcc vrf

To run on-demand route consistency checker (rcc) scan on AFI, SAFI, table, and prefix or the entire set of prefixes in the table, use the **show rcc vrf** command in XR EXEC mode.

show rcc { **ipv4** | **ipv6** } **unicast** *prefix / mask* **vrf** *vrfname*

Syntax Description

ipv4	IPv4 address prefix.
ipv6	IPv6 address prefix.
prefix /mask	Specifies unicast address prefix.
vrf	Specifies VPN routing and forwarding (VRF) instance.
vrfname	Name of the VRF.

Command Default

None.

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

Task ID

Task ID	Operation
ipv4	read
ipv6	read

This example shows how to run on-demand rcc scan for an IPv6 prefix:

```
RP/0/RP0/CPU0:router#show rcc ipv6 unicast 2001:DB8::/32 vrf vrf_1
```

This example shows how to run on-demand rcc scan for an Ipv4 prefix:

```
RP/0/RP0/CPU0:router#show rcc ipv4 unicast 10.2.3.4/32 vrf vrf-1
```

show rib

To display Routing Information Base (RIB) data, use the **show rib** command in XR EXEC mode.

show rib {**ipv4** | **ipv6**} {**unicast** | **multicast**}[**firsthop** | [*type interface-path-id*] | **next-hop** | [*type interface-path-id*] | **opaques** | {**attribute** | **ip-nexthop** | **ipfrr** | **safi-tunnel** | **summary** | **tunnel-nexthop**} | **protocols** | [**standby**] | **statistics** | [*name*] | [**standby**] | **topology** | {*topo-name* | **all**}]

Syntax Description

ipv4	(Optional) Specifies IP Version 4 address prefixes.
ipv6	(Optional) Specifies IP Version 6 address prefixes.
unicast	(Optional) Specifies unicast address prefixes. This is the default.
multicast	(Optional) Specifies multicast address prefixes.
firsthop	(Optional) Specifies registered first-hop notification addresses.
type	Interface type. For more information, use the question mark (?) online help function.
interface-path-id	Identifies a physical interface or a virtual interface. **Note** Use the show interfaces command to see a list of all possible interfaces currently configured on the router. For more information about the syntax for the router, use the question mark (?) online help function.
next-hop	(Optional) Specifies registered next-hop notification addresses.
opaques	(Optional) Specifies opaque data installed in the RIB.
attribute	(Optional) Specifies opaque attributes installed in the RIB.
ip-nexthop	(Optional) Specifies P next-hop data installed in the RIB.
ipfrr	(Optional) Specifies IP fast reroute (IPFRR) opaque data installed in the RIB.
safi-tunnel	(Optional) Specifies subaddress family (SAFI) tunnel opaque data installed in the RIB.
summary	(Optional) Specifies a summary of opaque data installed in the RIB.
tunnel-nexthop	(Optional) Specifies tunnel next-hop opaque data installed in the RIB.
protocols	(Optional) Specifies registered protocols.
statistics *name*	(Optional) Specifies RIB statistics of a given name.
standby	(Optional) Specifies standby information.
topology *topo-name*	(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 mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

Task ID

Task ID	Operations
ipv4	read

Examples

The following example illustrates the **show rib** command:

```
RP/0/RSP0RP0/CPU0:router# show rib

            ipv4 multicast


            topology
          BLUE

RP/0/RSP0RP0/CPU0:router# show rib topology BLUE ipv4 multicast protocols
Protocol   Handle   Instance
isis         0      mt
```

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**] [**multicast**] [**next-hop**] [**opaques**] [**protocols**] [**recursion-depth-max**] [**safi-all**] [**statistics**] [**tables**] [**trace**] [**unicast**] [**vpn-attributes**]

Syntax Description

attributes	(Optional) Displays all BGP attributes installed in RIB.
client-id	(Optional) Displays RIB client ID for longer history of redistributed routes sent to the client.
clients	(Optional) Displays RIB clients.
extcomms	(Optional) Displays all extended communities installed in RIB.
firsthop	(Optional) Displays registered firsthop notification addresses.
history	(Optional) Displays redistributed routes sent to RIB clients.
multicast	(Optional) Displays multicast commands.
next-hop	(Optional) Displays registered next-hop notification addresses.
opaques	(Optional) Displays opaquae data installed in RIB.
protocols	(Optional) Displays registered protocols.
recursion-depth-max	(Optional) Displays maximum recursion depth in RIB.
safi-all	(Optional) Displays unicast and multicast commands.
statistics	(Optional) Displays RIB statistics.
tables	(Optional) Displays a list of tables known to RIB.
trace	(Optional) Displays RIB trace entries.
unicast	(Optional) Displays unicast commands.
vpn-attributes	(Optional) Displays all VPN attributes installed in RIB.

Command Default

No default behavior or values

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

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.
```

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 mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

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

id	ID of the client. Range is 0 to 4294967295.
redistribution history	Displays longer history of redistributed routes sent to RIB clients.
standby	(Optional) Displays standby information.

Command Default

No default behavior or values

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

Use the **show rib client-id** command to display a history of the route additions, deletions, and updates sent from RIB to the client across VRFs.

Task ID

Task ID	Operations
rib	read

Examples

The following is sample output from the **show rib client-id** command:

```
RP/0/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 95: 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.

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

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 mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

Use the **show rib clients** command to display the list of clients who have registered with RIB, what protocol routes they are redistributing, and a history of the routes sent to the client.

The maximum number of redistribution entries is 5000 for Bulk Content Downloader (BCDL) and 500 for other protocols.

Task ID

Task ID	Operations
rib	read

Examples

The following is sample output from the **show rib clients** command:

```
RP/0/RP0/CPU0:router# show rib clients

Process             Location         Client ID  Redist   Proto
isis                node0_5_CPU0     0           insync   insync
ospf                node0_5_CPU0       2         insync   insync

RP/0/RP0/CPU0:router# show rib clients redistribution

isis node0_5_CPU0
```

```
   ipv4  uni        vrf default   insync         route
      static                      insync
ospf node0_5_CPU0
   ipv4  uni        vrf default   insync         route
      static                      insync
      local                       insync
bgp node0_5_CPU0
   ipv4  uni        vrf abc       insync         route
      static                      insync
bcdl_agent node0_5_CPU0
   ipv4  uni        vrf default   insync         rib_fib
   ipv4  uni        vrf bar       insync         rib_fib
   ipv4  uni        vrf abc       insync         rib_fib
   ipv4  uni        vrf test      insync         rib_fib
```

This table describes the significant fields shown in the display.

Table 96: show rib clients Field Descriptions

Field	Description
Process	Client process name.
Location	Location where the client process in running.
Client ID	ID assigned to the client by RIB.
Redist	Whether the client is redistributing any protocols or not and whether it has read all routes from RIB or not. • insync—read • outsync—not read.
Proto	Whether the protocol has sent all its routes to RIB and signaled update complete or not. • insync—read • outsync—not read.

show rib extcomms

To display all extended communities installed in the Routing Information Base (RIB), use the **show rib extcomms** command in 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

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

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
```

This table describes the significant fields shown in the display.

Table 97: show rib extcomms Field Descriptions

Field	Description
Extended Community	Type of extended communities. Different protocols can add different extended communities.

Field	Description
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 .

show rib [**vrf** {*vrf-name* | **all**}] [**afi-all** | **ipv4** | **ipv6**] [**unicast** | **multicast** | **safi-all**] **firsthop** [*client-name*] [*type interface-path-id* | *ip-address* /*prefix-length* | *ip-address mask* | **resolved** | **unresolved** | **damped**] [**summary**] [**standby**]

Syntax Description

vrf { *vrf-name* | **all** }	(Optional) Specifies a particular VPN routing and forwarding (VRF) instance or all VRF instances.
afi-all	(Optional) Specifies all address families.
ipv4	(Optional) Specifies IP Version 4 address prefixes. This is the default.
ipv6	(Optional) Specifies IP Version 6 address prefixes.
unicast	(Optional) Specifies unicast address prefixes. This is the default.
multicast	(Optional) Specifies multicast address prefixes.
safi-all	(Optional) Specifies unicast and multicast address prefixes.
client-name	(Optional) Name of the RIB client.
type	Interface type. For more information, use the question mark (**?**) online help function.
interface-path-id	Physical interface or virtual interface. **Note** Use the **show interfaces** command to see a list of all interfaces currently configured on the router. For more information about the syntax for the router, use the question mark (**?**) online help function.
ip-address	(Optional) Network that BGP advertises.
/ *prefix-length*	(Optional) Length of the IP address prefix. A decimal value that indicates how many of the high-order contiguous bits of the address compose the prefix (the network portion of the address). A slash (/) must precede the decimal value.
ip-address mask	(Optional) Network mask applied to the *ip-address* argument.
resolved	(Optional) Specifies resolved next-hops.
unresolved	(Optional) Specifies unresolved next-hops.
damped	(Optional) Specifies next-hops that are damped.
summary	(Optional) Specifies a summary of the next-hop information.
standby	(Optional) Displays standby information.

Command Default

If a **vrf** *vrf-name* is not specified, routes are cleared from the default IPv4 unicast VRF.

Command Modes

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

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

afi-all	(Optional) Specifies all address families.
ipv4	(Optional) Specifies IP Version 4 address prefixes. This is the default.
ipv6	(Optional) Specifies IP Version 6 address prefixes.
client-id *client-id*	(Optional) Specifies the ID of the client. Range for *client-id* argument is 0 to 4294967295.
standby	(Optional) Displays standby information.

Command Default

No default behavior or values

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

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 98: 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 [**vrf** {*vrf-name* | **all**}] [**afi-all** | **ipv4** | **ipv6**] [**unicast** | **multicast** | **safi-all**] **next-hop** [*client-name*] [*type interface-path-id* | *ip-address /prefix-length* | *ip-address mask* | **resolved** | **unresolved** | **damped**] [**summary**] [**standby**]

Syntax Description

vrf { *vrf-name* | **all** }	(Optional) Specifies a particular VPN routing and forwarding (VRF) instance or all VRF instances.
afi-all	(Optional) Specifies all address families.
ipv4	(Optional) Specifies IP Version 4 address prefixes. This is the default.
ipv6	(Optional) Specifies IP Version 6 address prefixes.
unicast	(Optional) Specifies unicast address prefixes. This is the default.
multicast	(Optional) Specifies multicast address prefixes.
safi-all	(Optional) Specifies unicast and multicast address prefixes.
client-name	(Optional) Name of the RIB client.
type	Interface type. For more information, use the question mark (**?**) online help function.
interface-path-id	Physical interface or virtual interface. **Note** Use the **show interfaces** command to see a list of all interfaces currently configured on the router. For more information about the syntax for the router, use the question mark (**?**) online help function.
ip-address	(Optional) Network IP address about which routing information should be displayed.
mask	(Optional) Network mask specified in either of two ways: • Network mask can be a four-part, dotted-decimal address. For example, 255.0.0.0 indicates that each bit equal to 1 means the corresponding address bit is a network address. • Network mask can be indicated as a slash (/) and number. For example, /8 indicates that the first 8 bits of the mask are 1s, and the corresponding bits of the address are the network address.
/ *prefix-length*	(Optional) Length of the IP address prefix. A decimal value that indicates how many of the high-order contiguous bits of the address compose the prefix (the network portion of the address). A slash (/) must precede the decimal value.

resolved	(Optional) Specifies resolved next-hops.
unresolved	(Optional) Specifies unresolved next-hops.
damped	(Optional) Specifies next-hops that are damped.
summary	(Optional) Specifies a summary of the next-hop information.
standby	(Optional) Displays standby information.

Command Default

No default behavior or values

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 24.4.1	This command output is extended to include • new values (`recurse cross table` and `connected interfaces only`) in the `Flags` field and • new field `metric type`.
Release 7.0.12	This command was introduced.

Usage Guidelines

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
```

This example shows the output for **show rib next-hop** command with the new values in the `Flags` field and the new field `metric type`:

```
Router #show rib next-hop 203.0.113.100/32
Firsthop prefix: 203.0.113.100/32
 Flags: recurse, recurse cross table, connected interfaces only
 Last event occurred Nov 16 00:04:28.940, 00:00:51 ago; version 2

 Registered clients:
```

```
rib_lib_test/node0_0_CPU0 created Nov 16 00:04:28.939, 00:00:51 ago
read last notification at Nov 16 00:04:28.943, 00:00:51 ago
reference count is 1
Flags: pack MLDP root

Destination paths:
192.0.2.1 - Bundle-Ether100
198.51.100.2 - Bundle-Ether101
203.0.113.3 - Bundle-Ether102 in vrf test
192.0.2.4 - Bundle-Ether103 in vrf test

Resolving route: 203.0.113.100/32 known via "bgp rib_lib_test_xtf"
Metric computed: 0
Metric type: aigp of resolving route
```

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 [**vrf** {*vrf-name* | **all**}] [**afi-all** | **ipv4** | **ipv6**] [**unicast** | **multicast** | **safi-all**] **opaques** {**attribute** | **ip-nexthop** | **ipfrr** | **safi-tunnel** | **summary** | **tunnel-nexthop**} [*rib-client-name*] [**standby**]

Syntax Description

vrf { *vrf-name* | **all** }	(Optional) Specifies a particular VPN routing and forwarding (VRF) instance or all VRF instances.
afi-all	(Optional) Specifies all address families.
ipv4	(Optional) Specifies IP Version 4 address prefixes. This is the default.
ipv6	(Optional) Specifies IP Version 6 address prefixes.
unicast	(Optional) Specifies unicast address prefixes. This is the default.
multicast	(Optional) Specifies multicast address prefixes.
safi-all	(Optional) Specifies unicast and multicast address prefixes.
attribute	Displays opaque attributes installed in the RIB.
ip-nexthop	Displays IP next-hop data installed in the RIB.
ipfrr	Displays IP fast reroute (IPFRR) opaque data installed in the RIB. **Note** Since the IP/LDP per-prefix LFA-FRR feature was introduced in IOS XR Software Release 4.0.1, the **show rib opaques ipfrr** command has been deprecated. Use **show route** command as part of the per-prefix LFA-FRR feature to determine backup paths. • **show route** output displays all FRR Backup paths. The FRR Backup paths are indicated with a (!). • **show route detail** output displays path ID and backup-path ID to identify if a path is protected and if so by which path.
safi-tunnel	Displays subaddress family (SAFI) tunnel opaque data installed in the RIB.
summary	Displays a summary of opaque data installed in the RIB.
tunnel-nexthop	Displays tunnel next-hop opaque data installed in the RIB.
rib-client-name	(Optional) Name of the RIB client.
standby	(Optional) Displays standby information.

Command Default

No default behavior or values

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

If information is not used by the RIB server process, it is viewed as opaque data. Use the **show rib opaques** command to display opaque data installed in the RIB.

Task ID

Task ID	Operations
rib	read

Examples

The following is sample output from the **show rib opaques** command:

```
  RP/0/RP0/CPU0:router# show rib opaques safi-tunnel

  Summary of safi tunnel opaque data in IPv4 RIB:

  Opaque key: 1:10.1.0.2
  Opaque data:
  Tunnel Encap - ifhandle=0x1000180, type=L2TPv3, Params=[Session-id=0x1EB1127C, `
Cookielen=8, Cookie=0xA73A3E0AFCD419A6] Opaque key: 65535:10.0.101.1 Opaque data:


  RP/0/RP0/CPU0:router# show rib ipv6 opaques tunnel-nexthop

  Summary of 6PE/6VPE IP over tunnel nexthop opaque data in IPv6 RIB:

  Opaque key: 1:::ffff:10.1.0.2
  Opaque key: 65535:::ffff:10.0.101.1
  Opaque key: 65535:::ffff:10.0.101.2
  Opaque key: 65535:::ffff:10.0.101.3
  Opaque key: 65535:::ffff:10.0.101.4
  Opaque key: 65535:::ffff:10.0.101.5
```

This table describes the significant fields shown in the display.

Table 99: show rib opaques Field Descriptions

Field	Description
Opaque key	Unique key for the opaque data as populated by the protocol client.
Opaque data	Data for the given key.

show rib protocols

To display protocols registered for route addition, use the **show rib protocols** command in XR EXEC mode.

show rib [**vrf** {*vrf-name* | **all**}] [**afi-all** | **ipv4** | **ipv6**] [**unicast** | **multicast** | **safi-all**] **protocols** [**standby**]

Syntax Description

vrf { *vrf-name* | **all** }	(Optional) Specifies a particular VPN routing and forwarding (VRF) instance or all VRF instances.
afi-all	(Optional) Specifies all address families.
ipv4	(Optional) Specifies IP Version 4 address prefixes. This is the default.
ipv6	(Optional) Specifies IP Version 6 address prefixes.
unicast	(Optional) Specifies unicast address prefixes. This is the default.
multicast	(Optional) Specifies multicast address prefixes.
safi-all	(Optional) Specifies unicast and multicast address prefixes.
standby	(Optional) Displays standby information.

Command Default

If a **vrf** *vrf-name* is not specified, the registered first-hop notification addresses are displayed for the default IPv4 unicast VRF.

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

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 100: 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 mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

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 101: 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 [**vrf** {*vrf-name* | **all**}] [**afi-all** | **ipv4** | **ipv6**] [**unicast** | **multicast** | **safi-all**] **statistics** [*client-name*] [**standby**]

Syntax Description

vrf { *vrf-name* | **all** }	(Optional) Specifies a particular VPN routing and forwarding (VRF) instance or all VRF instances.
afi-all	(Optional) Specifies all address families.
ipv4	(Optional) Specifies IP Version 4 address prefixes. This is the default.
ipv6	(Optional) Specifies IP Version 6 address prefixes.
unicast	(Optional) Specifies unicast address prefixes. This is the default.
multicast	(Optional) Specifies multicast address prefixes.
safi-all	(Optional) Specifies unicast and multicast address prefixes.
client-name	(Optional) Name of the RIB client.
standby	(Optional) Displays standby information.

Command Default

If a **vrf** *vrf-name* is not specified, routes are cleared from the default IPv4 unicast VRF.

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

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_SUCESS at Jun 26 21:43:33.601

 Received 217422 light weight messages
  4 route add requests, 2 route delete requests
  10 protocol registered, 1 protocol unregistered
  0 protocol modify, 0 protocol purged
  14 protocol redistributions, 0 unregistered protocol redistributions
  0 reset protocol redistributions
  3 first hop registered, 1 first hop unregistered
  3 advertisements, 0 unregistered advertisement
  57 bind data, 97 update completes, 217230 other requests
   udp/node0_RP0_CPU0 last performed firsthop lookup operation
    with status success at Jun 27 10:09:59.990

 Received 0 nexthop batch messages
    0 successes
    0 inits
    0 registers, 0 unregisters
    0 register complete, 0 sync unregistered, 0 batch finished
```

This table describes the significant fields shown in the display.

Table 102: show rib statistics Field Descriptions

Field	Description
Received	Statistics received including batch messages and route, attribute, complete, and convergent operations.
Results of the batch message received	Batch message results.
Received *n* light weight messages	Number of lightweight API messages sent from RIB clients.
Received *n* nexthop batch messages	Number of batch API messages sent from RIB clients received by the RIB.

show rib tables

To display all tables known to the Routing Information Base (RIB), use the **show rib tables** command in 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 mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

Use the **show rib tables** command to display all tables known to the RIB, including table attributes. Attributes include VPN routing and forwarding (VRF) instance, address family, and maximum prefix information.

Task ID

Task ID	Operations
rib	read

Examples

The following is sample output from the **show rib tables** command when entered without an address:

```
RP/0/RP0/CPU0:router# show rib tables

Codes: N - Prefix Limit Notified, F - Forward Referenced
       D - Table Deleted, C - Table Reached Convergence

VRF             SAFI  Table ID     PrfxLmt   PrfxCnt TblVersion  N F D C
default         uni   0xe0000000   2000000        72        137  N N N Y
default         multi 0xe0100000   2000000         0          0  N N N Y
```

This table describes the significant fields shown in the display.

Table 103: show rib tables Field Descriptions

Field	Description
VRF	Name of the VRF instance.
SAFI	Subaddress family instance.
Table ID	ID of the RIB table.
PrfxLmt	Configured prefix limit for the RIB table.
PrfxCnt	Number of configured prefixes in the RIB table.
TblVersion	Tables version number.
N	Message sent when prefix limit is exceeded.
F	Forward referenced. If Y is indicated, a table has been created by RIB because a client has registered for the table, but RIB has not heard from the router space infrastructure (RSI) about the table. RSI manages the tables.
D	If Y is indicated, the table has been deleted in the RSI but RIB has not cleared the information.
C	Table reached convergence.

show rib trace

To display all Routing Information Base (RIB) library call tracer (ltrace) entries, use the **show rib trace** command in 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 *node-id* argument is entered in the *rack/slot/module* notation.
location { **all** | *node-id* }	(Optional) Displays ltrace entries for the designated node. The *node-id* argument is entered in the *rack/slot/module* notation. The **all** keyword displays ltrace entries for all nodes.

Command Default

No default behavior or values

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

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/RSP0RP0/CPU0 t1  Create: Management thread

  Mar 16 14:59:27.959 rib/ipv4_rib/rib-startup 0/RSP0RP0/CPU0 t2  Create: Management event
 manager
  Mar 16 14:59:28.346 rib/ipv4_rib/rib-io 0/RSP0RP0/CPU0 t1  Initialise: RIB server
  Mar 16 14:59:28.346 rib/ipv4_rib/rib-io 0/RSP0RP0/CPU0 t1  Initialise: Client collection
  Mar 16 14:59:28.676 rib/ipv4_rib/rib-io 0/RSP0RP0/CPU0 t1  Initialise: DB collection
  Mar 16 14:59:28.693 rib/ipv4_rib/rib-io 0/RSP0RP0/CPU0 t1  Initialise: Timer tree
  Mar 16 14:59:28.694 rib/ipv4_rib/rib-io 0/RSP0RP0/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/RSP0RP0/CPU0 t2  Initialise: Debugging
routine
  Mar 16 14:59:29.128 rib/ipv4_rib/rib-io 0/RSP0RP0/CPU0 t1  Register: read, select cb
functions
  Mar 16 14:59:29.137 rib/ipv4_rib/rib-startup 0/RSP0RP0/CPU0 t1  Register: cerrno DLL name
 lib_rib_error.dll
  .
  .
  .
```

show rib vpn-attributes

To display all VPN attributes installed in the Routing Information Base (RIB), use the **show rib vpn-attributes** command in XR EXEC mode.

show rib [**afi-all** | **ipv4** | **ipv6**] **vpn-attributes** [**summary**] [**standby**]

Syntax Description

afi-all	(Optional) Specifies all address families.
ipv4	(Optional) Specifies IP Version 4 address prefixes.
ipv6	(Optional) Specifies IP Version 6 address prefixes.
summary	(Optional) Displays VPN attribute information.
standby	(Optional) Displays standby information.

Command Default

The default is IPv4 address prefixes.

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

Task ID

Task ID	Operations
rib	read

Examples

The following is sample output from the **show rib vpn-attributes** command:

```
RP/0/RP0/CPU0:router# show rib vpn-attributes

Extended community data in RIB:

Extended community                                Ref count
COST:128:128:41984                                        2
COST:128:129:42240                                        2
COST:128:129:44544                                        1
COST:128:129:169984                                       2
COST:128:129:307200                                       1

MVPN attribute data in RIB:

MVPN Attribute                                        Ref count
0:0:1:f4:0:0:0:1:1:1:1:1                                      1
0:0:2:bc:0:0:0:1:3:3:3:3                                     10
```

```
0:0:2:bc:0:0:0:1:3:3:3:4                                2
```

This table describes the significant fields shown in the display.

Table 104: show rib vpn-attributes Field Descriptions

Field	Description
Extended Community	Extended community added by the protocol clients.
Ref Count	Number of routes referring to the same extended community.
MVPN Attribute	Connector attribute added by BGP to support MVPNs.
Ref Count	Number of routes referring to the same extended community.

show rib vrf

To display all VRF table information in the Routing Information Base (RIB), use the **show rib vrf** command in XR EXEC mode.

show rib vrf {*vrf-name* | **all**} [**ipv4**] [**ipv6**] [**afi-all**] [**firsthop**] [**next-hop**] [**opaques**] [**protocols**] [**statistics** *name*]

Syntax Description

vrf { *vrf-name* | **all** }	(Optional) Specifies a particular VPN routing and forwarding (VRF) instance or all VRF instances.
ipv4	(Optional) Specifies IP Version 4 address prefixes.
ipv6	(Optional) Specifies IP Version 6 address prefixes.
afi-all	(Optional) Specifies all address families.
firsthop	(Optional) Specifies registered first-hop notification addresses
next-hop	(Optional) Specifies registered next-hop notification addresses.
opaques	(Optional) Specifies opaque data installed in the RIB.
protocols	(Optional) Specifies registered protocols.
statistics *name*	(Optional) Specifies RIB statistics for the given name.

Command Default

No default behavior or values

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

Task ID

Task ID	Operations
ipv4	read

Examples

The following example shows output from the **show rib vrf all statistics** command:

```
RP/0/RP0/CPU0:router# show rib vrf all statistics
RP/0/RP0/CPU0:router#
```

show route

To display the current routes in the Routing Information Base (RIB), use the **show route** command in XR EXEC mode.

show route [**vrf** {*vrf-name* | **all**}] [**afi-all** | **ipv4** | **ipv6**] [**unicast** | **multicast topology** *topo-name* | **safi-all**] [**protocol** [*instance*] | *ip-address* [*mask*] | *ip-address*/*prefix-length*] [**standby**] [**detail**]

Syntax Description

vrf { *vrf-name* | **all** }	(Optional) Specifies a particular VPN routing and forwarding (VRF) instance or all VRF instances.
afi-all	(Optional) Specifies all address families.
ipv4	(Optional) Specifies IP Version 4 address prefixes. This is the default.
ipv6	(Optional) Specifies IP Version 6 address prefixes.
unicast	(Optional) Specifies unicast address prefixes. This is the default.
multicast	(Optional) Specifies multicast address prefixes.
topology *topo-name*	(Optional) Specifies topology table information and name of the topology table.
safi-all	(Optional) Specifies unicastand multicast address prefixes.
protocol	(Optional) Name of a routing protocol. If you specify a routing protocol, use one of the following keywords: • **bgp** • **isis** • **ospf** • **rip** • **static** • **local** • **connected**
instance	(Optional) Number or name used to identify an instance of the specified protocol.
ip-address	(Optional) Network IP address about which routing information should be displayed.
mask	(Optional) Network mask specified in either of two ways: • Network mask can be a four-part, dotted-decimal address. For example, 255.0.0.0 indicates that each bit equal to 1 means the corresponding address bit is a network address. • Network mask can be indicated as a slash (/) and number. For example, /8 indicates that the first 8 bits of the mask are 1s, and the corresponding bits of the address are the network address.
/prefix-length	(Optional) Length of the IP address prefix. A decimal value that indicates how many of the high-order contiguous bits of the address compose the prefix (the network portion of the address). A slash (/) must precede the decimal value.

standby	(Optional) Displays standby information.
detail	(Optional) Displays detailed information for the specified prefix.

Command Default

If a **vrf** *vrf-name* is not specified, routes are cleared from the default IPv4 unicast VRF.

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

When the **afi-all** keyword is used, the *ip-address* and *mask* arguments are not available.

The **topology** keyword must be accompanied by the **ipv4 multicast** keywords, except when the **afi-all** keyword or the **safi-all** keyword is specified.

Task ID

Task ID	Operations
rib	read

Examples

The following is sample output from the **show route** command when entered without an address:

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

Codes: C - connected, S - static, R - RIP, B - BGP
       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, HundredGigE 0/3/0/0
L    3.2.3.2/32 is directly connected, 00:04:39, HundredGigE 0/3/0/0
O E2 5.2.5.0/24 [110/20] via 3.3.3.1, 00:04:20, HundredGigE 0/3/0/0
O E2 6.2.6.0/24 [110/20] via 3.3.3.1, 00:04:20, HundredGigE 0/3/0/0
C    7.2.7.0/24 is directly connected, 00:04:20, HundredGigE 0/3/0/7
L    7.2.7.2/32 is directly connected, 00:04:20, HundredGigE 0/3/0/7
O E2 8.2.8.0/24 [110/20] via 3.3.3.1, 00:04:20, HundredGigE 0/3/0/0

C    10.3.0.0/16 is directly connected, 13:14:59, HundredGigE 0/0/0/0
L    10.3.0.2/32 is directly connected, 13:14:59, HundredGigE 0/0/0/0
```

This table describes the significant fields shown in the display.

Table 105: show route Field Descriptions

Field	Description
S*	Code indicating how the route was derived. See the code legend preceding the output. In this case, the route was derived from a static (candidate default).
[1/0]	First number in the brackets is the administrative distance of the information source; the second number is the metric for the route.
1.0.0.0/16	Address and prefix length of the remote network.
MgmtEthernet 0/5/CPU0/0	Specifies the interface through which the specified network can be reached.
C	Code indicating how the route was derived. See the code legend preceding the output. In this case, the route was connected.
L	Code indicating how the route was derived. See the code legend preceding the output. In this case, the route was local.
O	Code indicating how the route was derived. See the code legend preceding the output. In this case, the route was on-demand routing (ODR).
E2	Code indicating how the route was derived. See the code legend preceding the output. In this case, the route was OSPF external type 2.
8.2.8.0/24	Address and prefix length of the remote network connected to the static route.
via 3.3.3.1	Specifies the address of the next router to the remote network.
13:14:59	Specifies the last time the route was updated.
(!)	Code indicating fast re-route (FRR) backup path information.

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 HundredGigE 0/0/0/0
       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 HundredGigE 2/1
       Route metric is 10
    No advertising protos.
```

This table describes the significant fields shown in the display.

Table 106: show route with IP Address Field Descriptions

Field	Description
Routing entry for	Network address and mask.
Known via	Indicates how the route was derived.
distance	Administrative distance of the information source.
metric	Route value assigned by the routing protocol.
type	IS-IS type level.
Routing Descriptor Blocks:	Displays the next-hop IP address followed by the information source.
from ... via ...	First address is the next-hop IP address, and the other is the information source. This report is followed by the interface for this route.
Route metric	Best metric for this Routing Descriptor Block.

No advertising protos.	Indicates that no other protocols are advertising the route to their redistribution consumers. If the route is being advertised, protocols are listed in the following manner:
	
	`Redist Advertisers:`
	`  isis p`
	`  ospf 43`

The following example illustrates the **show route** command with the **topology** *topo-name* keyword and argument specified:

```
RP/0/RP0/CPU0:router# show route ipv4 multicast topology green

  Codes: C - connected, S - static, R - RIP, B - BGP
```

```
       O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP
       i - ISIS, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, su - IS-IS summary null, * - candidate default
       U - per-user static route, o - ODR, L - local, G  - DAGR
       A - access/subscriber, (!) - FRR Backup path

  Gateway of last resort is not set

  i L1 10.1.102.0/24 [115/20] via 10.1.102.41, 1w4d, HundredGigE 0/1/0/0.1
  i L1 10.3.3.0/24 [115/20] via 10.1.102.41, 1w4d, HundredGigE 0/1/0/0.1
  i L1 192.168.0.40/32 [115/20] via 10.1.102.41, 1w4d, HundredGigE 0/1/0/0.1
```

This example is a sample **show route summary** command output that displays fast-reroute (FRR) Backup path information. The FRR Backup paths are indicated with a (!).

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

Codes: C - connected, S - static, R - RIP, B - BGP
       O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP
       i - ISIS, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, su - IS-IS summary null, * - candidate default
       U - per-user static route, o - ODR, L - local, G  - DAGR
       A - access/subscriber, (!) - FRR Backup path

Gateway of last resort is not set

B    1.2.3.4/32 [200/0] via 10.10.1.3, 00:01:40
C    2.0.0.0/30 is directly connected, 03:28:47, ServiceApp40
L    2.0.0.1/32 is directly connected, 03:28:47, ServiceApp40
C    2.0.1.0/30 is directly connected, 03:13:05, ServiceApp43
L    2.0.1.1/32 is directly connected, 03:13:05, ServiceApp43
C    2.4.1.0/24 is directly connected, 03:11:35, HundredGigE 0/4/0/0
L    2.4.1.2/32 is directly connected, 03:11:35, HundredGigE 0/4/0/0
C    3.1.0.0/30 is directly connected, 03:33:48, ServiceInfra1
L    3.1.0.2/32 is directly connected, 03:33:48, ServiceInfra1
C    3.1.3.0/30 is directly connected, 03:18:14, ServiceInfra2
L    3.1.3.2/32 is directly connected, 03:18:14, ServiceInfra2
C    5.3.0.0/16 is directly connected, 03:58:29, MgmtEth0/RP0/CPU0/0
                is directly connected, 03:58:29, MgmtEth0/RP1/CPU0/0
L    5.3.16.10/32 is directly connected, 03:59:07, MgmtEth0/RP1/CPU0/0
L    5.3.16.12/32 [0/0] via 5.3.16.12, 03:58:29, MgmtEth0/RP0/CPU0/0
L    5.3.16.16/32 is directly connected, 03:58:29, MgmtEth0/RP0/CPU0/0
B    5.4.0.0/16 [200/0] via 10.1.1.10, 00:01:36
S    5.10.0.0/16 [1/0] via 5.3.0.1, 03:59:07
O    10.1.1.3/32 [110/11] via 40.1.10.1, 00:00:17, Bundle-Ether10
                 [110/11] via 200.40.1.101, 00:00:17, Bundle-Ether1.1
                 [110/0] via 100.100.2.1, 00:00:17, HundredGigE 0/2/0/3.1 (!)
L    10.1.1.6/32 is directly connected, 03:58:29, Loopback0
O    10.1.1.9/32 [110/22] via 40.1.10.1, 00:00:17, Bundle-Ether10
                 [110/22] via 200.40.1.101, 00:00:17, Bundle-Ether1.1
                 [110/0] via 100.100.2.1, 00:00:17, HundredGigE 0/2/0/3.1 (!)
O    10.1.1.10/32 [110/111] via 40.1.10.1, 00:00:17, Bundle-Ether10
                  [110/111] via 200.40.1.101, 00:00:17, Bundle-Ether1.1
                  [110/0] via 100.100.2.1, 00:00:17, HundredGigE 0/2/0/3.1 (!)
O    10.1.1.11/32 [110/0] via 40.1.1.1, 00:01:33, Bundle-Ether1 (!)
                  [110/101] via 40.3.3.2, 00:01:33, HundredGigE 0/5/0/9
O    10.1.1.12/32 [110/111] via 40.1.10.1, 00:00:17, Bundle-Ether10
```

```
                  [110/111] via 200.40.1.101, 00:00:17, Bundle-Ether1.1
                  [110/0] via 100.100.2.1, 00:00:17, HundredGigE 0/2/0/3.1 (!)
O    10.1.1.16/32 [110/21] via 40.1.10.1, 00:00:17, Bundle-Ether10
                  [110/21] via 200.40.1.101, 00:00:17, Bundle-Ether1.1
                  [110/0] via 100.100.2.1, 00:00:17, HundredGigE 0/2/0/3.1 (!)
```

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 HundredGigE 0/2/0/3.1, 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.
```

This example is a sample **show route ipv6** command output:

```
RP/0/RP0/CPU0:router#show route ipv6
Fri May 18 14:00:10.996 EDT
Codes: C - connected, S - static, R - RIP, B - BGP
O - OSPF, IA - OSPF inter area
N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP
i - ISIS, L1 - IS-IS level-1, L2 - IS-IS level-2
ia - IS-IS inter area, su - IS-IS summary null, * - candidate default
U - per-user static route, o - ODR, L - local, G - DAGR
A - access/subscriber, (!) - FRR Backup path
Gateway of last resort is not set
C 1111:2222::abcd/128 is directly connected,
06:20:02, HundredGigE 0/0/0/4
```

This example is a sample **show route ipv6 detail** command output:

```
RP/0/RP0/CPU0:router#show route ipv6 1111:2222::abcd/128 detail
Fri May 18 14:00:20.798 EDT
```

```
Routing entry for 1111:2222::abcd/128
Known via "connected l2tpv3_xconnect", distance 0, metric 0 (connected)
Installed May 18 07:40:08.522 for 06:20:12
Routing Descriptor Blocks
1111:2222::abcd directly connected, via HundredGigE 0/0/0/4
Route metric is 0
Label: 0x2 (2)
Tunnel ID: None
Extended communities count: 0
Route version is 0xd (13)
No local label
IP Precedence: Not Set
QoS Group ID: Not Set
Route Priority: RIB_PRIORITY_CONNECTED (2) SVD Type RIB_SVD_TYPE_LOCAL
Download Priority 0, Download Version 13
No advertising protos.
```

This example is a sample **show route ipv6 summary** command output:

```
RP/0/RP0/CPU0:router#show route ipv6 summary
Fri May 18 14:00:28.988 EDT
Route Source Routes Backup Deleted Memory (bytes)
local 0 0 0 0
connected l2tpv3_xconnect 1 0 0 160
connected 0 0 0 0
Total 1 0 0 160
```

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 [**vrf** {*vrf-name* | **all**}] [**afi-all** | **ipv4** | **ipv6**] [**unicast** | **multicast** | {**topology** *topo-name*} | **safi-all**] **backup** [*ip-address [mask] ip-address /prefix-length]*][**standby**]

Syntax Description

vrf { *vrf-name* | **all** }	(Optional) Specifies a particular VPN routing and forwarding (VRF) instance or all VRF instances.
afi-all	(Optional) Specifies all address families.
ipv4	(Optional) Specifies IP Version 4 address prefixes.
ipv6	(Optional) Specifies IP Version 6 address prefixes.
unicast	(Optional) Specifies unicast address prefixes.
multicast	(Optional) Specifies multicast address prefixes.
topology *topo-name*	(Optional) Specifies topology table information and name of the topology table.
safi-allsafi-all	(Optional) Specifies unicast and multicast address prefixes.
ip-address	(Optional) Network IP address about which backup routing information should be displayed.
mask	(Optional) Network mask specified in either of two ways: • Network mask can be a four-part, dotted decimal address. For example, 255.0.0.0 indicates that each bit equal to 1 means the corresponding address bit is a network address. • Network mask can be indicated as a slash (/) and number. For example, /8 indicates that the first 8 bits of the mask are ones, and the corresponding bits of the address are the network address.
/prefix-length	(Optional) Length of the IP address prefix. A decimal value that indicates how many of the high-order contiguous bits of the address compose the prefix (the network portion of the address). A slash (/) must precede the decimal value.
standby	(Optional) Displays standby information.

Command Default

If a **vrf** *vrf-name* is not specified, routes are cleared from the default IPv4 unicast VRF.

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

Use the **show route backup** command to display information about routes that have been installed into the RIB as backup routes. This command also displays information about the currently selected active route for which there is a backup.

When the **afi-all** keyword is used, the *ip-address* and *mask* arguments are not available.

The **topology** keyword must be accompanied by the **ipv4 multicast** keywords, except when the **afi-all** keyword or the **safi-all** keyword is specified.

Task ID

Task ID	Operations
rib	read

Examples

The following is sample output from the **show route backup** command:

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

Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
       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, HundredGigE 4/0/0/1
                  Backup  O E2 [110/1] via 10.12.12.2, HundredGigE 3/0/0/1
```

This table describes the significant fields shown in the display.

Table 107: 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.
HundredGigE4/0/0/1	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.
HundredGigE 3/0/0/1	Outbound interface for the OSPF version of the route.

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 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.
ip-address	IP address about which best local information should be displayed.
standby	(Optional) Displays standby information.

Command Modes

Usage Guidelines

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:

```
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 108: show route best-local Field Descriptions

Field	Description
Routing entry for	Identifies the requested IP address.
Known via	Indicates how the route was derived.

Field	Description
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	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 [**vrf** {*vrf-name* | **all**}] [**afi-all** | **ipv4** | **ipv6**] [**unicast** | **multicast** | {**topology** *topo-name*} | **safi-all**] **connected** [**standby**]

Syntax Description

vrf { *vrf-name* | **all** }	(Optional) Specifies a particular VPN routing and forwarding (VRF) instance or all VRF instances.
afi-all	(Optional) Specifies all address families.
ipv4	(Optional) Specifies IP Version 4 address prefixes.
ipv6	(Optional) Specifies IP Version 6 address prefixes.
unicast	(Optional) Specifies unicast address prefixes.
multicast	(Optional) Specifies multicast address prefixes.
topology *topo-name*	(Optional) Specifies topology table information and name of the topology table.
safi-all	(Optional) Specifies unicast and multicast address prefixes.
standby	(Optional) Displays standby information.

Command Default

If a **vrf** *vrf-name* is not specified, routes are cleared from the default IPv4 unicast VRF.

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

Use the **show route connected** command to display information about connected routes in the routing table.

The topology keyword must be accompanied by the ipv4 multicast keywords, except when the afi-all keyword or the safi-all keyword is specified.

Task ID

Task ID	Operations
rib	read

Examples

The following is sample output from the **show route 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, HundredGigE 0/3/0/0
C    7.7.7.0/24 is directly connected, 00:33:00, HundredGigE 0/3/0/7
C    10.0.0.0/16 is directly connected, 13:43:40, HundredGigE 0/0/0/0
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 109: 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.

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 [**vrf** {*vrf-name* | **all**}] [**afi-all** | **ipv4** | **ipv6**] [**unicast** | **multicast** | {**topology** *topo-name*} | **safi-all**] **local** [*type interface -path-id*] [**standby**]

Syntax Description

vrf { *vrf-name* | **all** }	(Optional) Specifies a particular VPN routing and forwarding (VRF) instance or all VRF instances.
afi-all	(Optional) Specifies all address families.
ipv4	(Optional) Specifies IP Version 4 address prefixes.
ipv6	(Optional) Specifies IP Version 6 address prefixes.
unicast	(Optional) Specifies unicast address prefixes.
multicast	(Optional) Specifies multicast address prefixes.
topology *topo-name*	(Optional) Specifies topology table information and name of the topology table.
safi-all	(Optional) Specifies unicast and multicast address prefixes.
type	Interface type. For more information, use the question mark (**?**) online help function.
interface-path-id	Physical interface or virtual interface. **Note** Use the **show interfaces** command to see a list of all interfaces currently configured on the router. For more information about the syntax for the router, use the question mark (**?**) online help function.
standby	(Optional) Displays standby information.

Command Default

If a **vrf** *vrf-name* is not specified, routes are cleared from the default IPv4 unicast VRF.

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

Use the **show route local** command to display information about local routes in the routing table.

The topology keyword must be accompanied by the ipv4 multicast keywords, except when the afi-all keyword or the safi-all keyword is specified.

Task ID

Task ID	Operations
rib	read

Examples

The following is sample output from the **show route local** command:

```
RP/0/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, HundredGigE6/0/0/1
L    172.22.12.1/32 is directly connected, 00:13:35, HundredGigE3/0/0/1
L    192.168.20.2/32 is directly connected, 00:13:27, HundredGigE4/0/0/1
L    10.254.254.1/32 is directly connected, 00:13:26, HundredGigE5/0/0/1
```

This table describes the significant fields shown in the display.

Table 110: 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.

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 [**vrf** {*vrf-name* | **all**}] [**ipv4** | **ipv6**] [**unicast** | **multicast** | {**topology** *topo-name*} | **safi-all**] **longer-prefixes** {*ip-address mask ip-address/prefix-length*} [**standby**]

Syntax Description

vrf { *vrf-name* | **all** }	(Optional) Specifies a particular VPN routing and forwarding (VRF) instance or all VRF instances.
ipv4	(Optional) Specifies IP Version 4 address prefixes.
ipv6	(Optional) Specifies IP Version 6 address prefixes.
unicast	(Optional) Specifies unicast address prefixes.
multicast	(Optional) Specifies multicast address prefixes.
topology *topo-name*	(Optional) Specifies topology table information and name of the topology table.
safi-all	(Optional) Specifies unicast and multicast address prefixes.
ip-address	Network IP address about which routing information should be displayed.
mask	Network mask specified in either of two ways: • Network mask can be a four-part, dotted-decimal address. For example, 255.0.0.0 indicates that each bit equal to 1 means the corresponding address bit is a network address. • Network mask can be indicated as a slash (/) and number. For example, /8 indicates that the first 8 bits of the mask are 1s, and the corresponding bits of the address are the network address.
/ *prefix-length*	Length of the IP address prefix. A decimal value that indicates how many of the high-order contiguous bits of the address compose the prefix (the network portion of the address). A slash (/) must precede the decimal value.
standby	(Optional) Displays standby information.

Command Default

If a **vrf** *vrf-name* is not specified, routes are cleared from the default IPv4 unicast VRF.

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

Use the **show route longer-prefixes** command to troubleshoot forwarding problems whose cause may be a long prefix.

The topology keyword must be accompanied by the ipv4 multicast keywords, except when the afi-all keyword or the safi-all keyword is specified.

Task ID

Task ID	Operations
rib	read

Examples

The following is sample output from the **show route longer-prefixes** command:

```
RP/0/RP0/CPU0:router# show route longer-prefixes 172.16.0.0/8

Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
       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/RSP0RP0/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/RSP0RP0/CPU0/0
```

This table describes the significant fields shown in the display.

Table 111: 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 *nd*n*h*) since the route was installed in the RIB.
MgmtEth0/RSP0 RP0/CPU0/0	Outbound interface for the route.

show route next-hop

To filter routes by the next-hop address or interface, use the **show route next-hop** command in 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.
ip-address	(Optional) IP address about which next-hop information is to be displayed.
standby	(Optional) Displays standby information.

Command Modes

Usage Guidelines

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:

```
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:

```
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 112: 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 *ndn*h) since the route was installed in the RIB.
GigabitEthernet0/1/0/2	Outbound interface for the route.

Related Commands

Command	Description
Show route	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 [**vrf** {*vrf-name* | **all**}] [**ipv4** | **ipv6**] [**unicast** | **multicast** | {**topology** *topo-name*} | **safi-all**] **quarantined** [*ip-address/prefix-length]ip-address mask*] [**standby**]

Syntax Description

vrf { *vrf-name* | **all** }	(Optional) Specifies a particular VPN routing and forwarding (VRF) instance or all VRF instances.
ipv4	(Optional) Specifies IP Version 4 address prefixes.
ipv6	(Optional) Specifies IP Version 6 address prefixes.
unicast	(Optional) Specifies unicast address prefixes.
multicast	(Optional) Specifies multicast address prefixes.
topology *topo-name*	(Optional) Specifies topology table information and name of the topology table.
safi-all	(Optional) Specifies unicast and multicast address prefixes.
ip-address	(Optional) IP address about which looping routes information is to be displayed.
/ *prefix-length*	(Optional) Length of the IP address prefix. A decimal value that indicates how many of the high-order contiguous bits of the address compose the prefix (the network portion of the address). A slash (/) must precede the decimal value.
ip-address mask	(Optional) Network mask applied to the *ip-address* argument.
standby	(Optional) Displays standby information.

Command Default

If a **vrf** *vrf-name* is not specified, routes are cleared from the default IPv4 unicast VRF.

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

RIB quarantining detects mutually recursive routes and quarantines the last route that actually completes the mutual recursion. The quarantined route is periodically evaluated to see if the mutual recursion has gone away. If the recursion still exists, the route remains quarantined. If the recursion has gone away, the route is released from quarantine.

Use the **show route quarantined** command to display mutually recursive (looping) routes.

The topology keyword must be accompanied by the ipv4 multicast keywords, except when the afi-all keyword or the safi-all keyword is specified.

Task ID

Task ID	Operations
rib	read

Examples

The following is sample output from the **show route quarantined** command:

```
RP/0/RP0/CPU0:routerr# show route quarantined

Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
       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 113: show route quarantined Field Descriptions

Field	Description
10.10.109.1/32	IP address and length of the route.
[1/0]	Distance and metric for the route.
via 10.10.34.1	IP address of next-hop on the route.
00:00:01	Time (in hh:mm:ss or *n*d*n*h) since the route was installed in the RIB.
(quarantined)	Shows that the route is quarantined.

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 [**vrf** {*vrf-name* | **all**}] [**ipv4** | **ipv6**] [**unicast** | **multicast** | {**topology** *topo-name*} | **safi-all**] **resolving-next-hop** *ip-address* [**standby**]

Syntax Description

vrf { *vrf-name* | **all** }	(Optional) Specifies a particular VPN routing and forwarding (VRF) instance or all VRF instances.
ipv4	(Optional) Specifies IP Version 4 address prefixes.
ipv6	(Optional) Specifies IP Version 6 address prefixes.
unicast	(Optional) Specifies unicast address prefixes.
multicast	(Optional) Specifies multicast address prefixes.
topology *topo-name*	(Optional) Specifies topology table information and name of the topology table.
safi-all	(Optional) Specifies unicast and multicast address prefixes.
ip-address	IP address about which resolved next-hop information is to be displayed.
standby	(Optional) Displays standby information.

Command Default

If a **vrf** *vrf-name* is not specified, routes are cleared from the default IPv4 unicast VRF.

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

Use the **show route resolving-next-hop** command to perform a recursive route lookup on the supplied destination address and return information on the next immediate router (next hop) to the destination.

The topology keyword must be accompanied by the ipv4 multicast keywords, except when the afi-all keyword or the safi-all keyword is specified.

Task ID

Task ID	Operations
rib	read

Examples

The following is sample output from the **show route resolving-next-hop** command:

```
RP/0/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 114: 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.

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 [**vrf** {*vrf-name* | **all**}] [**afi-all** | **ipv4** | **ipv6**] [**unicast** | **multicast** | {**topology** *topo-name*} | **safi-all**] **static** [**standby**]

Syntax Description

vrf { *vrf-name* | **all** }	(Optional) Specifies a particular VPN routing and forwarding (VRF) instance or all VRF instances.
afi-all	(Optional) Specifies all address families.
ipv4	(Optional) Specifies IP Version 4 address prefixes.
ipv6	(Optional) Specifies IP Version 6 address prefixes.
unicast	(Optional) Specifies unicast address prefixes.
multicast	(Optional) Specifies multicast address prefixes.
topology *topo-name*	(Optional) Specifies topology table information and name of the topology table.
safi-all	(Optional) Specifies unicast and multicast address prefixes.
standby	(Optional) Displays standby information.

Command Default

If a **vrf** *vrf-name* is not specified, routes are cleared from the default IPv4 unicast VRF.

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

Use the **show route static** command to display information about static routes in the routing table.

The topology keyword must be accompanied by the ipv4 multicast keywords, except when the afi-all keyword or the safi-all keyword is specified.

Task ID

Task ID	Operations
rib	read

Examples

The following is sample output from the **show route static** command:

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

```
S    10.1.1.0/24 is directly connected, 00:54:05, HundredGigE3/0/0/1
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 115: 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.
HundredGigE3/0/0/1	Outbound interface for the route.
[1/0]	Distance and metric for the route.

show route summary

To display the current contents of the Routing Information Base (RIB), use the **show route summary** command in XR EXEC mode mode.

show route [**vrf** {*vrf-name* | **all**}] [**afi-all** | **ipv4** | **ipv6**] [**unicast** | **multicast** | {**topology** *topo-name*} | **safi-all**] **summary** [**detail**] [**standby**]

Syntax Description

vrf { *vrf-name* | **all** }	(Optional) Specifies a particular VPN routing and forwarding (VRF) instance or all VRF instances.
afi-all	(Optional) Specifies all address families.
ipv4	(Optional) Specifies IP Version 4 address prefixes.
ipv6	(Optional) Specifies IP Version 6 address prefixes.
unicast	(Optional) Specifies unicast address prefixes.
multicast	(Optional) Specifies multicast address prefixes.
topology *topo-name*	(Optional) Specifies topology table information and name of the topology table.
safi-all	(Optional) Specifies unicast and multicast address prefixes.
detail	(Optional) Displays a detailed summary of the contents of the RIB, including the number of paths and some protocol-specific route attributes.
standby	(Optional) Displays standby information.

Command Default

If a **vrf** *vrf-name* is not specified, routes are cleared from the default IPv4 unicast VRF.

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

Use the **show route summary** command to display information about routes in the routing information base.

When a route summary is needed frequently—for instance, in a polling situation—use the **show route summary** command without the **detail** keyword. The **detail** keyword is used less frequently for verification purposes, because it is much more expensive (in bandwidth), requiring a scan of the entire routing database.

The topology keyword must be accompanied by the ipv4 multicast keywords, except when the afi-all keyword or the safi-all keyword is specified.

Task ID

Task ID	Operations
rib	read

Examples

This example provides **show route summary** command output.

```
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 116: 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.

This example provides the output summary for all VRF instances.

```
RP/0/RP0/CPU0:DJ-SF-R3#show route vrf all ipv4 summary

Wed Jul 17 23:57:29.912 UTC
VRF: apple
Route Source              Routes     Backup     Deleted     Memory(bytes)
local                     2          0          0           384
connected                 2          0          0           384
dagr                      0          0          0           0
bgp                       714        50038      0           9608096
Total                     50042      0          0           9608864

VRF: apple2
Route Source              Routes     Backup     Deleted     Memory(bytes)
connected                 2          0          0           384
local                     2          0          0           384
dagr                      0          0          0           0
bgp                       714        36         0           7712
Total                     40         0          0           8480

VRF: apple3
Route Source              Routes     Backup     Deleted     Memory(bytes)
connected                 2          0          0           384
local                     2          0          0           384
dagr                      0          0          0           0
bgp                       714        36         0           7712
Total                     40         0          0           8480
```

```
VRF: **iid
Route Source              Routes     Backup     Deleted     Memory(bytes)
local                     0          0          0           0
connected                 0          0          0           0
Total                     0          0          0           0
```

Here, **iid** is the internal ID created for VRF in the system. It is required for some routing features.

This example provides the **show route summary** command output 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 117: show route summary detail Field Descriptions

Field	Description
Route Source	Source of the route. Routing protocol name and type.
Active Route	Number of active routes present in the routing table for each route source.
Active Path	Number of active paths present in the routing table for each route source.
Backup Route	Number of routes that are backup to a selected route for each route source.
Backup Path	Number of paths that are backup to a selected path for each route source.

table ip-only activate vrf

To install IP-only routes in the new VRF instance, use the **table ip-only activate vrf** command in RIB configuration mode.

table ip-only activate vrf *vrf_name*

Syntax Description

vrf *vrf_name*	Specifies an IP-only VPN routing and forwarding (VRF) instance.

Command Default

IP-only routes are not enabled.

Command Modes

Router RIB Configuration mode

Command History

Release	Modification
Release 24.2.11	This command was introduced.

Usage Guidelines

You must create a new VRF instance to handle IP-only routing tables before executing the **table ip-only activate vrf** *vrf_name* command.

Task ID

Task ID	Operations
rib	read, write

Examples

The following example shows how to enable IS-IS to install IP-only routes in the new VRF:

```
Router# configure
Router(config)# router rib
Router(config-rib)# table ip-only activate vrf ip_only
```

tag-map tag

To assign or map specific tags to routes in the Routing Information Base (RIB) for both IPv4 and IPv6 unicast routes in the VRF instance, use the **tag-map tag** <*value*>**map forwarding-hierarchy level-2-used-as-nexthop** command in the router configuration mode.

tag-map tag *value* **map forwarding-hierarchy level-2-used-as-nexthop**

Syntax Description

<*value*>	Specify the required tag value to the routes you want to map. The range is 1 to 4294967295.
map forwarding-hierarchy	Set the flag value of the route you want to map.
level-2-used-as-nexthop	Set the routes with matching tag to the flag value.

Command Default

By default, tag-map is not used.

Command Modes

Router Configuration mode.

Command History

Release	Modification
Release 24.4.1	This command was introduced.

Usage Guidelines

None.

Task ID

Task ID	Operations
rib	read, write

Example

The following example shows how to assign or map specific tags to routes in the RIB for IPv4 and IPv6 unicast routes in the VRF instance. The RIB maps all the IPv4 and IPv6 routes tagged with the value 100, which indicates that these routes resolve as nexthop for Layer 3 prefixes.

```
Router#config
Router(config)#router rib
Router(config-rib)#tag-map tag 100 map forwarding-hierarchy level-2-used-as-nexthop
Router(config-rib)#commit
```



RPL Commands

- as-path in, on page 655
- as-path-set, on page 657
- community matches-any, on page 659
- community matches-every, on page 661
- community-set, on page 663
- destination in, on page 666
- done, on page 668
- drop, on page 670
- drop (ePBR), on page 671
- end-policy, on page 672
- end-set, on page 673
- if, on page 675
- local-preference, on page 680
- med, on page 681
- next-hop-type is, on page 682
- pass, on page 683
- prefix-set, on page 684
- rib-has-route, on page 686
- route-policy (RPL), on page 687
- set community, on page 689
- set extcommunity transitive-bandwidth, on page 691
- set local-preference, on page 692
- set med, on page 693
- set next-hop, on page 695
- set origin, on page 697
- show rpl, on page 698
- show rpl active as-path-set, on page 700
- show rpl active community-set, on page 702
- show rpl active extcommunity-set, on page 704
- show rpl active prefix-set, on page 707
- show rpl active rd-set, on page 709
- show rpl active route-policy, on page 711
- show rpl as-path-set, on page 713

- show rpl as-path-set attachpoints, on page 714
- show rpl as-path-set references, on page 716
- show rpl community-set, on page 718
- show rpl community-set attachpoints, on page 720
- show rpl community-set references, on page 722
- show rpl extcommunity-set, on page 724
- show rpl inactive as-path-set, on page 727
- show rpl inactive community-set, on page 729
- show rpl inactive extcommunity-set, on page 731
- show rpl inactive prefix-set, on page 733
- show rpl inactive rd-set, on page 735
- show rpl inactive route-policy, on page 737
- show rpl maximum, on page 739
- show rpl policy-global references, on page 741
- show rpl prefix-set, on page 743
- show rpl prefix-set attachpoints, on page 744
- show rpl prefix-set references, on page 746
- show rpl rd-set, on page 748
- show rpl rd-set attachpoints, on page 749
- show rpl rd-set references, on page 751
- show rpl route-policy, on page 753
- show rpl route-policy attachpoints, on page 756
- show rpl route-policy inline, on page 758
- show rpl route-policy references, on page 760
- show rpl route-policy uses, on page 763
- show rpl unused as-path-set, on page 766
- show rpl unused community-set, on page 769
- show rpl unused extcommunity-set, on page 772
- show rpl unused prefix-set, on page 773
- show rpl unused rd-set, on page 776
- show rpl unused route-policy, on page 777
- transmit (ePBR), on page 780

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

as-path-set-name	Name of an AS path set.
inline-as-path-set	Inline AS path set. The inline AS path set must be enclosed in parentheses.
parameter	Parameter name. The parameter name must be preceded with a "$."

Command Default

No default behavior or values

Command Modes

Route-policy configuration

Usage Guidelines

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:

```
Router(config)# as-path-set my-as-set
Router(config-as)# ios-regex '_12$',
Router(config-as)# ios-regex '_13$'
Router(config-as)# end-set
```

and the following policy excerpt using an *as-path-set-name* argument:

```
Router(config-rpl)# if as-path in my-as-set then
Router(config-rpl-if)# set local-preference 100
Router(config-rpl-if)# endif
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:

```
Router(config-rpl)# if as-path in (ios-regex '_12$,ios-regex '_13$') then
Router(config-rpl-if)# set local-preference 100
Router(config-rpl-if)# endif
Router(config-rpl)#
```

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*

Syntax Description

name	Name of the AS path set.

Command Default

No default behavior or values

Command Modes

XR Config mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

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.
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.

Options	Description
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)#
```

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-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

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

Use the **community matches-any** command as a conditional expression within an **if** statement to match any element of a community set.

Note For a list of all conditional expressions available within an **if** statement, see the **if** command.

A simple condition using the **matches-any** operator evaluates as true if at least one community element of the community attribute for the route matches an element in the community set operand. If no community in the route matches any of the specifications in the named or inline set, then the condition evaluates to false. Likewise, when there is no community at all in the route, the condition evaluates to false.

Matching of a community in the route to a specification in a named or an inline set is intuitive. If the community specification in a set is the familiar colon-separated decimal 16-bit numbers specification, or one of the well-known communities, the community matches the specification if the specification denotes the same 32-bit number as that in the route. If the community specification uses a wildcard, then the community in the route matches if it is one of the many communities denoted by the wildcard specification. In inline sets, community specifications may be parameterized, in which case the relevant matching is done when the value of the parameter has been supplied.

Communities may also be matched using range and regular expression operators. Range specifications are entered as follows: [*low-value* .. *high-value*]. Either or both colon-separated halves of a community value may contain a range. The following are valid range specifications:

```
10:[100..1000]
[10..100]:80
[10..100]:[100..2000]
```

In addition, the **private-as** keyword may be used to specify the range from 64512 to 65534. Regular expressions are specified as the **ios-regex** keyword followed by a valid regular expression string.

Community values from the route are matched one at a time to the match specifications. Therefore, regex match specifications are expected to represent one individual community value and not a sequence of community values.

Task ID

Task ID	Operations
route-policy	read, write

Examples

In the following example, a named community set called my-community-set and a route policy called community-matches-any-example are created. The policy sets the local-preference to 100 for any route that has one or more of the communities in the my-community-set community set. If the route does not have any of these communities, the policy checks whether it has any communities whose first half is in the range from 10 to 25 and whose second half is the value 35, in which case it sets the local-preference to 200. Otherwise, it checks for a community value in the range of 30:100 to 30:500, in which case it sets the local-preference to 300.

```
  RP/0/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
```

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-name 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

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

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-rp-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
```

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*

Syntax Description

name	Name of the community set.

Command Default

No default behavior or values

Command Modes

XR Config mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

Regular expressions and ranges can be specified to match the communities. An attempt to use a community set that contains a range or regular expression to set a community value is rejected when an attempt to attach such a policy is made.

A community set holds community values for matching against the Border Gateway Protocol (BGP) community attribute. A community is a 32-bit quantity. For notational convenience, each community value must be split in half and expressed as two unsigned decimal integers in the range from 0 to 65535, separated by a colon.

The inline form of a community set also supports parameterization. Each 16-bit portion of the community may be parameterized.

The routing policy language (RPL) provides symbolic names for the standard well-known community values: **accept-own** is 0xFFFF0001, **internet** is 0:0, **no-export** is 65535:65281, **no-advertise** is 65535:65282, and **local-as** is 65535:65283.

RPL also provides a facility for using wildcards in community specifications. A wildcard is specified by inserting an asterisk (*) in place of one of the 16-bit portions of the community specification, which indicates that any value for that portion of the community matches.

Every community set must contain at least one community value. An empty community set is invalid and the policy configuration system rejects it.

Community sets can be entered in these formats:

Format	Description
#-remark	Remark beginning with '#'
*****	Wildcard (any community or part thereof)
0-65535	16-bit half-community number
[	Left bracket to begin range
accept-own	Accept-Own (BGP well-known community)

Format	Description
dfa-regex	DFA (deterministic finite automata) style regular expression
internet	Internet (BGP well-known community)
ios-regex	Traditional IOS style regular expression
local-AS	Do not send outside local AS (BGP well-known community)
no-advertise	Do not advertise to any peer (BGP well-known community)
no-export	Do not export to next AS (BGP well-known community)
private-as	Match within BGP private AS range [64512..65534]

Note The dfa-regex and ios-regex syntax for community set is *"[']*[^':&<>]*:*[^':&<>]*[']"*. This means that regex starts with a single-quote (') followed by a string of any character (that does not include single-quote, colon, ampersand, less-than, greater-than, or space) followed by a colon, and a string of any characters (that does not include single-quote, colon, ampersand, less-than, greater-than, or space) followed by single-quote.

Task ID

Task ID	Operations
route-policy	read, write

Examples

In the following example, a community set named cset_accept_own is created:

```
RP/0/RP0/CPU0:router#configure
RP/0/RP0/CPU0:router(config)#community-set cset_accept_own
RP/0/RP0/CPU0:router(config-comm)#accept-own
RP/0/RP0/CPU0:router(config-comm)#end-set
```

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
```

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	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 7.0.12	This command was introduced.

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate taskIDs. 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.
- For a list of all conditional expressions available within an **if** statement, see the **if** command.
- This command takes either a named prefix set or an inline prefix set value as an argument. The condition returns true if the destination entry matches any entry in the prefix set or inline prefix set. An attempt to match a destination using a prefix set that is defined but contains no elements returns false.
- The routing policy language (RPL) provides the ability to test destinations for a match to a list of prefix match specifications using the **in** operator. The **destination in** command is protocol-independent.
- Use the **destination-prefix in** command to filter a flowspec route in BGP, and the **match destination-address** command to create a flowspec rule that matches a specific destination.
- Use the **destination in** command to filter BGP routes for other address families, and use the **destination-prefix in** command to filter flowspec routes in BGP that include a **match destination-address** clause.
- RPL supports both 32-bit IPv4 prefixes, specified in dotted-decimal format, and 128-bit IPv6 prefixes, specified in colon-separated hexadecimal format.

Task ID

Task ID	Operations
route-policy	read, write

Examples

In the following example, a prefix set named my-prefix-set is defined and a route policy named use-destination-in is created. Within the use-destination-in route policy, the **destination in** command is used within an **if** statement to learn if the destination is in the prefix-set named my-prefix-set. If it is, then local preference is set to 100. If it is not in my-prefix-set but does match the next prefix specifications, then local preference is set to 200.

```
  RP/0/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 7.0.12	This command was introduced.

Usage Guidelines

Use the **done** command to stop executing the policy and accept the route.

Note The **done** command can be used as an action statement within an **if** statement. For a list of all action statements available within an **if** statement, see the **if** command.

When encountering a **done** statement the route is passed and no further policy statements are executed. All modifications made to the route prior to the **done** statement are still valid.

Note The default action of a route policy is to drop or discard any routes that have not been either explicitly passed or for which no attempt has been made to modify with an action. The routing policy language (RPL) does not have specific "match clauses," which means the default drop behavior is controlled by whether a route has been explicitly passed or an attempt has been to modify the route using an action statement.

Task ID

Task ID	Operations
route-policy	read, write

Examples

In the following example, if the destination match succeeds for 29.0.0.0/8 le 32, the execution continues past set community 102:12 and onto the next statement. If the destination match succeeds for 39.0.0.0/8 le 32 execution, then the policy execution stops when in encounters the *done* statement.

```
RP/0/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 7.0.12	This command was introduced.

Usage Guidelines 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)#
```

drop (ePBR)

The command discards packets that match the associated class map criteria when configured under a PBR class within a policy map.

drop

Command Default

None.

Command Modes

Policy map configuration mode.

Command History

Release	Modification
Release 25.3.1	This command was introduced.

Task ID

Task ID	Operation
qos	read, write

```
Router(config)#policy-map type pbr INGRESS-SECURITY-POLICY
Router(config-pmap)#class type traffic UNWANTED-SOURCE-TRAFFIC
Router(config-pmap-c)#drop
Router(config-pmap-c)#exit
```

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 7.0.12	This command was introduced.

Usage Guidelines

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
```

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 7.0.12	This command was introduced.

Usage Guidelines

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
```

if

To decide which actions or dispositions should be taken for a given route, use the **if** command in route-policy configuration mode.

if *conditional-expression* **then** *action-statement* [*action-statement*] [**elseif** *conditional-expression* **then** *action-statement* [*action-statement*]] [**else** *action-statement* [*action-statement*]] **endif**

Syntax Description

conditional-expression	Expression to decide which actions or dispositions should be taken for the given route.
then	Executes an action statement if the **if** condition is true.
elseif	Strings together a sequence of tests.
else	Executes an action statement if the **if** condition is false.
endif	Ends the **if** statement.
action-statement	Sequence of operations that modify a route.

Command Default

No default behavior or values

Command Modes

Route-policy configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

The **if** command uses a conditional expression to decide which actions or dispositions should be taken for a given route. Table 1 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 2 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 118: Conditional Expressions

Command	Description
as-path in	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.

Command	Description
as-path is-local	Determines if the router (or another router within this autonomous system or confederation) originated the route.
as-path length	Performs a conditional check based on the length of the AS path.
as-path neighbor-is	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	Tests an AS path against the AS sequence beginning with the AS number that originated a route.
as-path passes-through	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	Performs specific checks based on the length of the AS path.
community is-empty	Learns if a route has community attributes associated with it.
community matches-any	Matches any element of a community set.
community matches-every	Matches every element of a community set.
destination in	Matches a destination entry in a named prefix set or inline prefix set.
extcommunity rt is-empty	Learns if a route has RT extended community attributes associated with it.
extcommunity rt matches-any	Matches elements of an RT extended community set.
extcommunity rt matches-every	Matches every element of an RT extended community set.
extcommunity rt matches-within	Matches at least one element of a Border Gateway Protocol (BGP) route target (RT) extended community set.
extcommunity soo is-empty	Learns if a route has SoO extended community attributes associated with it.
extcommunity soo matches-any	Matches elements of an SoO extended community set.
extcommunity soo matches-every	Matches every element of an SoO extended community set.
local-preference	Specifies BGP local-preference attribute
med	Compares the MED to an integer value or a parameterized value.
next-hop in	Compares the next-hop associated with the route to data contained in either a named or an inline prefix set.
orf prefix in	Matches a prefix in a prefix set or an inline prefix set.
origin is	Tests the value of the origin attribute.
path-type is	Tests the path type.

Command	Description
protocol	Checks if a protocol is installing the route.
rd in	Compares the RD associated with the route to data contained in either a named or an inline RD set.
rib-has-route	Checks if a route is in the RIB.
route-has-label	Checks if a route has a Multiprotocol Label Switching (MPLS) label.
route-type is	Compares route types when redistribution is being performed into BGP, OSPF, or IS-IS.
source in	Tests the source of the route against the data in either a named or an inline prefix set.
tag	Matches a specific tag value.
tag in	Conditionally compares tag-route against tag-set.
vpn-distinguisher is	Compares the VPN distinguisher against a specified value.

Table 119: Action Statements

Command	Description
abort (RPL)	Discards a route policy definition and returns to XR Config mode.
add	Adds an offset to an existing value.
apply	Executes a parameterized or an unparameterized policy from within another policy.
delete community	Deletes community values from a community list in a route.
delete extcommunity rt	Deletes extended community values from an extended community list in a route.
done	Accepts this route with no further processing
drop	Drops a route.
end-policy	Ends the definition of a route policy and exits route-policy configuration mode.
pass	Signifies that even though the route has not been modified, the user wants to continue executing in the policy block.
prepend as-path	Prepends the AS path with additional autonomous system numbers.
replace as-path	Replaces a sequence of AS numbers or private AS numbers in the AS path with the configured local AS.
set community	Sets the BGP community attribute.
set dampening	Configures BGP route dampening.
set extcommunity cost	Replaces or adds the extended communities for a cost on the route.

Command	Description
set extcommunity rt	Replaces or adds the extended communities for an RT on the route.
set ip-precedence	Sets the IP precedence to classify packets.
set isis-metric	Sets the IS-IS metric attribute value.
set label	Sets the BGP label attribute value.
set level	Configures the IS-IS level in which redistributed routes should be sent.
set local-preference	Specifies a preference value for the autonomous system path.
set med	Sets the MED value.
set metric-type (IS-IS)	Controls whether IS-IS treats the metric as an internal or external metric.
set metric-type (OSPF)	Controls whether OSPF treats the cost as a Type 1 or Type 2 metric.
set next-hop	Replaces the next-hop associated with a given route.
set origin	Changes the origin attribute.
set ospf-metric	Sets an OSPF protocol metric attribute value.
set qos-group (RPL)	Sets the QoS group to classify packets.
set rib-metric	Sets a RIB metric attribute value for a table policy.
set rip-metric	Sets RIP metric attributes.
set rip-tag	Sets route tag attribute.
set tag	Sets the tag attribute.
set traffic-index	Sets the traffic index attribute.
set weight	Sets the weight value for BGP routes.
suppress-route	Indicates that a given component of an aggregate should be suppressed, that is, not advertised.
unsuppress-route	Indicates that a given component of an aggregate should be unsuppressed.
set vpn-distinguisher	Sets the VPN distinguisher value.

Task ID

Task ID	Operations
route-policy	read, write

Examples

In the following example, any route whose AS path is in the set as-path-set-1 is dropped:

```
RP/0/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 Descriminator (MED) value of 150, then each route with both the community value of 12:34 or 56:78 and a MED of 150 is dropped.

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** } { *number parameter* }

Syntax Description

eq | **is** | **ge** | **le**	Equal to; exact match; greater than or equal to; less than or equal to.
number	Value assigned to a 32-bit unsigned integer. Range is 0 to 4294967295.
parameter	Parameter name. The parameter name must be preceded with a "$."

Command Default

No default behavior or values

Command Modes

Route-policy configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

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)#
```

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** } { *number parameter* }

Syntax Description

eq | **is** | **ge** | **le**	Equal to; exact match; greater than or equal to; less than or equal to.
number	Value assigned to a 32-bit unsigned integer. Range is 0 to 4294967295.
parameter	Parameter name. The parameter name must be preceded with a "$."

Command Default

No default behavior or values

Command Modes

Route-policy configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

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-type is

To filter routes based on Cisco Multiprotocol Label Switching Traffic Engineering tunnels as the next-hop type, use the **next-hop-type is** command in route-policy configuration mode.

next-hop-type is mpls-te

Syntax Description

mpls-te	Filters BGP routes based on Cisco Multiprotocol Label Switching Traffic Engineering tunnels.

Command Default

No default behavior or values.

Command Modes

Route-policy configuration

Command History

Release	Modification
Release 25.1.1	This command was introduced.

Usage Guidelines

None.

Task ID

Task ID	Operations
route-policy	read, write

Examples

The example configuration allows only Cisco Multiprotocol Label Switching Traffic Engineering routes by filtering nexthops based on the next-hop-type.

```
Router(config)#route-policy ROUTE-RESOLUTION
Router(config-rpl)#if protocol is isis 100 and route-type is level-1 then
Router(config-rpl-if)#pass
Router(config-rpl-if)#elseif protocol is isis 100 and route-type is level-2 and next-hop-type
 is mpls-te then
Router(config-rpl-elseif)#pass
Router(config-rpl-elseif)#else
Router(config-rpl-else)#drop
Router(config-rpl-else)#endif
Router(config-rpl)#end-policy
Router(config)#commit
```

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 7.0.12	This command was introduced.

Usage Guidelines

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)#
```

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*

Syntax Description

name	Name of a prefix set.

Command Default

None

Command Modes

XR Config mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

Use the **prefix-set** command to enter prefix set configuration mode and define a prefix set.

A prefix set is a comma-separated list of prefix match specifications. It holds IPv4 or IPv6 prefix match specifications, each of which has four parts: an address, a mask length, a minimum matching length, and a maximum matching length. The address is required, but the other three parts are optional. The address is a standard four-part, dotted-decimal numeric IPv4 address or a colon-separated hexadecimal IPv6 address. The mask length, if present, is a nonnegative decimal integer in the range from 0 to 32 for IPv4 prefixes or 0 to 128 for IPv6 prefixes following the address and separated from it by a slash. The optional minimum matching length follows the address and optional mask length and is expressed as the keyword **ge** (mnemonic for **g**reater than or **e**qual to), followed by a nonnegative decimal integer in the range from 0 to 32 for IPv4 or 0 to 128 for IPv6. The optional maximum matching length follows the rest and is expressed by the keyword **le** (mnemonic for **l**ess than or **e**qual to), followed by yet another nonnegative decimal integer in the range from 0 to 32 for IPv4 or 0 to 128 for IPv6. A syntactic shortcut for specifying an exact length for prefixes to match is the **eq** keyword, mnemonic for **eq**ual to.

If a prefix match specification has no mask length, then the default mask length is 32 for IPv4 or 128 for IPv6. The default minimum matching length is the mask length. If a minimum matching length is specified, then the default maximum matching length must be less than 32 for IPv4 prefixes or 128 for IPv6 prefixes. Otherwise, if neither a minimum nor maximum length is specified, the default maximum length is the mask length.

A prefix set is a list of prefix match specifications. It holds IPv4 or IPv6 prefix match specifications, each of which has two parts: an address and a mask. The address and mask is a standard dotted-decimal IPv4 or colon-separated hexadecimal IPv6 address. The prefix set allows the specifying of contiguous and non-contiguous set of bits that mus be matched in any route. The set of bits to be matched are provided in the form of a mask in which a binary 0 means a mandatory match and a binary 1 means a 'do not match' condition.

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
```

This example shows a prefix set named legal-ipv4-prefix:

```
RP/0/RP0/CPU0:router(config)# prefix-set legal-ipv4-prefix
RP/0/RP0/CPU0:router(config-pfx)# 10.1.1.1  0.255.0.255
RP/0/RP0/CPU0:router(config-pfx)# 10.2.2.2 0.0.0.0
RP/0/RP0/CPU0:router(config-pfx)# 10.3.3.3 255.255.255.254
RP/0/RP0/CPU0:router(config-pfx)# 10.4.4.4 255.255.255.255
```

In the above example, In the above example, the command defines the prefix-set named acl-prefix-set. The first element specifies to match all routes having 10 in first octet and 1 in third octet. The second element matches all routes having prefix as 10.2.2.2 (that is, matches all conditions). The third element matches all routes having odd numbers in the last octets and the fourth element matches all routes with any prefix.

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 7.0.12	This command was introduced.

Usage Guidelines

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-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)*]

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 mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

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.

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

community-set-name	Community set name.
inline-community-set	Inline community set. The inline community set must be enclosed in parentheses.
parameter	Parameter name. The parameter name must be preceded with a "$."
additive	(Optional) Adds communities to communities in the route.

Command Default

No default behavior or values

Command Modes

Route-policy configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

Use the **set community** command to set the BGP community attribute.

Note The **set community** command can be used as an action statement within an **if** statement. For a list of all action statements available within an **if** statement, see the **if** command.

Communities are 32-bit values carried in BGP routes. Each route may have zero or more communities in an unordered list.

Use this command to replace the communities in a route or add to them using the optional **additive** keyword.

As with the other community forms that support inline sets, either or both 16-bit portions of the community can be parameterized. Likewise, the names of the well-known communities **internet** (0:0), **no-advertise** (65535:65281), **no-export** (65535:65282), and **local-AS** (65535:65283) can also be used. In an inline community set, each 16-bit portion can also be specified as the **peeras** to express the AS number of the neighbor from which the route was received. If the neighbor AS employs a 4-byte ASN, the IANA-assigned 16-bit value 23456 (AS_TRANS) is used as **peeras** instead.

Without the **additive** keyword, any existing communities (other than the well-known communities) are removed and replaced with the given communities. The **additive** keyword specifies that all communities already present in the route be maintained and the list of communities be added to them.

Task ID

Task ID	Operations
route-policy	read, write

Examples

The following are incomplete configuration examples using the **set community** command:

```
RP/0/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 extcommunity transitive-bandwidth

To set and advertise the BGP DMZ external community attributes in a route with transitive-bandwidth extended community, use the **set extcommunity transitive-bandwidth** command in the route-policy configuration mode.

set extcommunity transitive-bandwidth { *parameter extcommunity-set-name* }

Syntax Description

extcommunity-set-name	External community set name.
parameter	Parameter name. You must precede the parameter name with a "$" when you want to use it as a variable.

Command Default

None.

Command Modes

Route-policy configuration.

Command History

Release	Modification
Release 25.1.1	The **transitive-bandwidth** type is introduced as an extended community in RPL.
Release 7.0.12	The **set extcommunity** command was introduced.

Task ID

Task ID	Operations
route-policy	read, write

Examples

The following is a configuration example using the **set extcommunity transitive-bandwidth** command, where you set and apply the transitive-bandwidth extended community for bandwidth advertisement in BGP:

```
Router#config
Router(config)#extcommunity-set transitive-bandwidth dmz-transitive
Router(config-ext)#1000:20000
Router(config-ext)#end-set
Router(config)#route-policy bgp_rpl_outbound
Router(config-rpl)#set extcommunity transitive-bandwidth dmz-transitive
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 { *number parameter* }

Syntax Description

number	Value assigned to a 32-bit unsigned integer. Range is 0 to 4294967295.
parameter	Parameter name. The parameter name must be preceded with a "$."

Command Default

Default value is 100.

Command Modes

Route-policy configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

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 8000 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 { *number parameter* | **igp-cost** | { **+** | { *number parameter* } | **-** | { *number parameter* } } | **max-reachable** }

Syntax Description

number	Value assigned to a 32-bit unsigned integer. Range is 0 to 4294967295.
parameter	Parameter name. The parameter name must be preceded with a "$."
igp-cost	Sets the MED value to the cost for the Interior Gateway Protocol (IGP) route to resolve the next-hop of the BGP route.
+ | -	Sets the MED to the MED plus or minus a static offset. An integer or parameter must follow the plus or minus.
max-reachable	Sets the MED value to the maximum possible value of 4294967295.

Command Default

No default behavior or values

Command Modes

Route-policy configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

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 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** *parameter* | **peer-address** | **self** | **unchanged** }

Syntax Description

ipv4-address	Valid IPv4 address.
ipv6-address	Valid IPv6 address.
discard	Sets next-hop as Null0 interface.
destination-vrf	(Optional) Specifies that the next-hop of the route should be resolved in destination VRF context. This keyword is available when an IPv4 or IPv6 address or parameter is used.
peer-address	Sets the next-hop to the IP address of the remote Border Gateway Protocol (BGP) peer.
parameter	Parameter name. The parameter name must be preceded with a "$."
self	Sets itself as the next-hop.
unchanged	Preserves the received next-hop without modification.

Command Default

No default behavior or values

Command Modes

Route-policy configuration

Command History

Release	Modification
Release 7.0.12	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 **destination-vrf** keyword is used mainly in Layer 3 VPN networks when importing routes.

The below address families support the selective setting of 'next-hop-self' via the RPL statement 'set next-hop self' starting in 4.2.1. Previous to this the setting of next-hop-self via an RPL was for all prefixes in the address family or none of the prefixes.

- IPv4 unicast
- IPv4 labeled-unicast
- IPv4 multicast
- IPv6 unicast
- IPv6 multicast

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.

Use the **unchanged** keyword when you want the route-policy to leave the next-hop attribute exactly as it was received, without making any changes. This is useful in scenarios where retaining the original next-hop is required for proper routing or policy consistency.

Task ID

Task ID	Operations
route-policy	read, write

Examples

In the following example, the next-hop is set to a valid IPv4 address:

```
Router(config-rpl)# set next-hop 10.0.0.5
```

In this example, the next-hop is set to a parameter value $nexthop:

```
Router(config-rpl)# set next-hop $nexthop
```

In this example, the next-hop is set to a valid IPv4 address with a destination VRF context:

```
Router(config-rpl)# set next-hop 10.0.0.5 destination-vrf
```

In the following example, the next-hop attribute is preserved as received by using the unchanged keyword in a route-policy:

```
Router(config-rpl)# set next-hop unchanged
```

set origin

To change the Border Gateway Protocol (BGP) origin attribute, use the **set origin** command in route-policy configuration mode.

set origin { **igp** | **incomplete** | **egp** *parameter* }

Syntax Description

igp	Sets the origin type to Interior Gateway Protocol (IGP).
incomplete	Sets the origin type to incomplete.
egp	Sets the origin type to Exterior Gateway Protocol (EGP).
parameter	Parameter name. The parameter name must be preceded with a "$."

Command Default

No default behavior or values

Command Modes

Route-policy configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

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
```

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 *configuration-limit*	(Optional) Displays the number of lines to which configuration is limited. Range is 1 to 131072. The *configuration-limit* argument is available if the **running-config** keyword is specified.
policies *policies-limit*	(Optional) Displays the limit on the number of policies. Range is 1 to 5000. The *configuration-limit* argument is available if the **running-config** keyword is specified.
editor	(Optional) Specifies the default RPL editor. This keyword is available if the **running-config** keyword is specified.
emacs	(Optional) Displays the default RPL editor to Micro Emacs.
nano	(Optional) Displays the default RPL editor to nano.
vim	(Optional) Displays the default RPL editor to Vim.

Command Default

No default behavior or values

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.
Release 7.11.1	The **Nano** and **Emacs** keyword was deprecated.

Usage Guidelines

No specific guidelines impact the use of this command.

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
!
```

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 mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

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
```

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 mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

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,
  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
```

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 mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

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:

```
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
```

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 mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

Use the **show rpl active prefix-set** command to display all prefix sets that are in use in the system and referenced either directly or indirectly at a policy attach point.

Task ID

Task ID	Operations
route-policy	read

Examples

This example shows the following sample configuration:

```
router bgp 2
 address-family ipv4 unicast
 !
 neighbor 10.0.101.2
  remote-as 100
  address-family ipv4 unicast
   route-policy policy_1 in
  !
 !
 neighbor 10.0.101.3
  remote-as 12
  address-family ipv4 unicast
   route-policy policy_2 in
  !
 !
!

RP/0/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
!
```

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
```

show rpl active rd-set

To display the route distinguisher (RD) sets that are referenced by at least one policy that is being used at an attach point, use the **show rpl active rd-set** command in EXEC mode.

show rpl active rd-set [**detail**]

Syntax Description

detail	(Optional) Displays the content of the object and all referenced objects for active route policies.

Command Default

No default behavior or values

Command Modes

EXEC mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

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
```

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 mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

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
```

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 mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

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
```

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

name	Name of an AS path set.

Command Default

No default behavior or values

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

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 120: 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.

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

name	Name of the prefix set.
brief	(Optional) Limits the output to just the brief table and not the detailed information for the named AS path set.

Command Default

No default behavior or values

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

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
  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 121: 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.

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 mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

Use the optional **brief** keyword to limit the display to a list of the names of community sets without their configurations.

Task ID

Task ID	Operations
route-policy	read

The following is the sample output of the show rpl community-set command with graceful maintenance feature attributes displayed:

```
RP/0/0/CPU0:R5#show rpl community-set
Thu Jan 29 17:55:04.792 PST
Listing for all Community Set objects

community-set gshut
  graceful-shutdown
end-set
```

Examples

This example shows the following sample configuration:

```
route-policy policy_4
  if (destination in prefix_set_ex2) then
    if (community matches-any comm_set_ex2) then
      set community  (10:666) additive
    endif
    if (extcommunity matches-any ext_comm_set_rt_ex2) then
      set community  (10:999) additive
    endif
  endif
```

```
end-policy
```

Given this sample configuration, the **show rpl community-set comm_set_ex2** command displays the following information:

```
RP/0/RP0/CPU0:router# show rpl community-set comm_set_ex2

community-set comm_set_ex2
  65501:1,
  65501:2,
  65501:3
end-set
```

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

name	Name of a community set.

Command Default

No default behavior or values

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

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 122: 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.

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

name	Name of a community set.
brief	(Optional) Limits the output to just the summary table and not the detailed information for the community set.

Command Default

No default behavior or values

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

Use the **show rpl community-set references** command to display all the policies that reference the named community set.

Use the optional **brief** keyword to limit the output to just a summary table and not the detailed information for the community set.

Task ID

Task ID	Operations
route-policy	read

Examples

This example shows the following sample configuration:

```
router bgp 2
 address-family ipv4 unicast
 !
 neighbor 10.0.101.3
  remote-as 12
  address-family ipv4 unicast
   route-policy policy_2 in
  !
 !
!
route-policy policy_2
  if (destination in prefix_set_ex1) then
    if (community matches-any comm_set_ex1) then
      set community  (10:666) additive
    endif
    if (extcommunity matches-any ext_comm_set_rt_ex1) then
      set community  (10:999) additive
    endif
  endif
```

```
end-policy
```

Given this sample configuration, the **show rpl extcommunity-set comm_set_ex1 references** command displays the following information:

```
RP/0/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 123: 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.

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

name	(Optional) Name of the community set.
attachpoints	(Optional) Displays all attach points for this community set.
references	(Optional) Displays all policies that use this community set.
cost	(Optional) Displays all extended community cost sets.
rt	(Optional) Displays all extended community RT sets.
soo	(Optional) Displays all extended community SoO sets.
brief	(Optional) Limits the display to a list of the names of all extended community sets without their configurations.
states	(Optional) Displays all unused, inactive, and active states.

Command Default

If an attachpoint or reference is not specified, all configured extended community sets are displayed

If a cost, RT, or SoO sets is not specified, all configured extended community sets are displayed

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

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 extsoo1
  66:70001
end-set
!
extcommunity-set rt rtsetl1
  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,
  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
!
```

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 mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

Use the **show rpl inactive as-path-set** command to display all AS path sets that are not in use at an attach point either directly or indirectly but are referenced by at least one policy in the system.

Task ID

Task ID	Operations
route-policy	read

Examples

This example shows the following sample configuration:

```
router bgp 2
 address-family ipv4 unicast
 !
 neighbor 10.0.101.2
  remote-as 100
  address-family ipv4 unicast
   route-policy policy_1 in
  !
 !
 neighbor 10.0.101.3
  remote-as 12
  address-family ipv4 unicast
   route-policy policy_2 in
  !
 !
!
route-policy sample
  if (destination in sample) then
    drop
  endif
end-policy
!
route-policy policy_1
  if (destination in prefix_set_ex1) then
    set local-preference 100
  endif
```

```
  if (as-path in as_path_set_ex1) then
    set community  (10:333) additive
  endif
end-policy
!
route-policy policy_2
   if destination in prefix_set_ex1 then
     if (community matches-any comm_set_ex1) then
       set community (10:666) additive
     endif
     if (extcommunity rt matches-any ext_comm_set_rt_ex1) then
       set community (10:999) additive
     endif
   endif
end-policy
!
route-policy policy_3
  if (destination in prefix_set_ex2) then
    set local-preference 100
  endif
  if (as-path in as_path_set_ex2) then
    set community  (10:333) additive
  endif
end-policy
!
route-policy policy_4
  if (destination in prefix_set_ex2) then
    if (community matches-any comm_set_ex2) then
      set community  (10:666) additive
    endif
    if (extcommunity matches-any ext_comm_set_rt_ex2) then
      set community  (10:999) additive
    endif
  endif
end-policy
!
route-policy policy_5
  apply sample1
  apply policy_3
end-policy
```

Given this sample configuration, the **show rpl inactive as-path-set** command displays the following information:

```
RP/0/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
```

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 mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

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
```

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 mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

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
```

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 mode

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

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
```

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 mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

Use the **show rpl inactive rd-set** command to display all RD sets that are not in use at an attach point either directly or indirectly but are referenced by at least one policy in the system.

Task ID

Task ID	Operations
route-policy	read

Examples

This example shows the following sample configuration:

```
rd-set rdset1
  10:151,
  100.100.100.1:153,
  100.100.100.62/31:63
end-set
!
rd-set rdset2
  10:152,
  100.100.100.1:154,
  100.100.100.62/31:89
end-set
!
```

Given this sample configuration, the **show rpl inactive rd-set** command displays the following information:

```
RP/0/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
```

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 mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

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
```

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 mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

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
```

show rpl maximum field descriptions describes the significant fields shown in the display.

Table 124: 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.

Field	Description
Compiled policies size (kB)	Displays the current compiled total for policies in kilobytes.

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 mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

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 mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

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
```

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

name	Name of a prefix set.

Command Default

No default behavior or values

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

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 125: 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.

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

name	Name of the prefix set.
brief	(Optional) Limits the output to just a summary table and not the detailed information for the named prefix set.

Command Default

No default behavior or values

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

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 126: 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.

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 mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

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
```

show rpl rd-set attachpoints

To display all the policies used at an attach point that reference the named route distinguisher (RD) set, use the **show rpl rd-set attachpoints** command in XR EXEC mode.

show rpl rd-set *name* **attachpoints**

Syntax Description

name	Name of an RD set.

Command Default

No default behavior or values

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

Use the **show rpl rd-set attachpoints** command to display all the policies used at an attach point that reference the named RD set either directly or indirectly.

Task ID

Task ID	Operations
route-policy	read

Examples

This example shows the following sample configuration:

```
route-policy rdsetmatch
  if rd in rdset1 then
    set community (10:112)
  elseif rd in rdset2 then
    set community (10:223)
  endif
end-policy

router bgp 10
address-family vpnv4 unicast
 exit
 neighbor 10.0.101.1
  remote-as 11
  address-family vpnv4 unicast
   route-policy rdsetmatch in
!
```

Given this sample configuration, the **show rpl rd-set rdset1 attachpoints** command displays the following information:

```
RP/0/RP0/CPU0:router# show rpl rd-set rdset attachpoints

BGP Attachpoint: Neighbor

Neighbor/Group  type  afi/safi   in/out   vrf name
```

```
------------------------------------------------
10.0.101.1      --    IPv4/vpn   in         default
```

This table describes the significant fields shown in the display.

Table 127: show rpl rd-set attachpoints Field Descriptions

Field	Description
Neighbor/Group	BGP neighbor or neighbor group where the specified RD is used.
afi/safi	BGP address family or subaddress family where the RD set is used.
in/out	Direction
vrf name	VRF name where the RD set is used.

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

name	Name of the RD set.
brief	(Optional) Limits the output to just a summary table and not the detailed information for the RD set.

Command Default

No default behavior or values

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

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 128: 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.

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 mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

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
  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
```

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

name	Name of a policy.

Command Default

No default behavior or values

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

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
!
```

The following command displays the route policy attach points for policy_2:

```
RP/0/RP0/CPU0:router# show rpl route-policy policy_2 attachpoints

BGP Attachpoint: Neighbor

Neighbor/Group  type  afi/safi   in/out   vrf name
-------------------------------------------------
10.0.101.2      --    IPv4/uni   in         default
10.0.101.2      --    IPv4/uni   out        default
```

This table describes the significant fields shown in the display.

Table 129: show rpl route-policy attachpoints Field Descriptions

Field	Description
BGP Attachpoint	Location of the attach point.
Neighbor/Group	IP address of the attach point on the neighbor.
type	Displays the address family mode.
afi/safi	Address family identifier or subsequent address family identifier.
vrf name	Name of the VPN routing and forwarding (VRF) instance.

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

name	Name of a policy.

Command Default

No default behavior or values

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

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

name	Name of a prefix set.
brief	(Optional) Limits the output to just a summary table and not the detailed information for the named policy.

Command Default

No default behavior or values

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

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 130: 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.

Field	Description
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.

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

name	Name of a policy.
policies	Generates a list of all policies that the named policy uses.
sets	Lists all named sets that are used by the policy.
all	Generates a list of both sets and policies that the named policy references.
direct	(Optional) Lists only the policies or sets used directly in the named policy block. Set or policy references that occur as a result of an **apply** statement are not listed.

Command Default

No default behavior or values

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

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 131: 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 mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

Use the **show rpl unused as-path-set** command to display all AS path sets that are not used in a policy at an attach point either directly or indirectly and are not referenced by any policies in the system.

Task ID

Task ID	Operations
route-policy	read

Examples

This example shows the following sample configuration:

```
router bgp 2
 address-family ipv4 unicast
 !
 neighbor 10.0.101.2
  remote-as 100
  address-family ipv4 unicast
   route-policy policy_1 in
  !
 !
 neighbor 10.0.101.3
  remote-as 12
  address-family ipv4 unicast
   route-policy policy_2 in
  !
 !
!
as-path-set as_path_set_ex1
  ios-regex '^_65500_$',
  ios-regex '^_65501_$'
end-set
!
as-path-set as_path_set_ex2
  ios-regex '^_65502_$',
  ios-regex '^_65503_$'
end-set
!
```

```
as-path-set as_path_set_ex3
  ios-regex '^_65504_$',
  ios-regex '^_65505_$'
end-set
!
route-policy sample
  if (destination in sample) then
    drop
  endif
end-policy
!
route-policy policy_1
  if (destination in prefix_set_ex1) then
    set local-preference 100
  endif
  if (as-path in as_path_set_ex1) then
    set community  (10:333) additive
  endif
end-policy
!
route-policy policy_2
  if (destination in prefix_set_ex1) then
    if (community matches-any comm_set_ex1) then
      set community  (10:666) additive
    endif
    if (extcommunity matches-any ext_comm_set_rt_ex1) then
      set community  (10:999) additive
    endif
  endif
end-policy
!
route-policy policy_3
  if (destination in prefix_set_ex2) then
    set local-preference 100
  endif
  if (as-path in as_path_set_ex2) then
    set community  (10:333) additive
  endif
end-policy
!
route-policy policy_4
  if (destination in prefix_set_ex2) then
    if (community matches-any comm_set_ex2) then
      set community  (10:666) additive
    endif
    if (extcommunity matches-any ext_comm_set_rt_ex2) then
      set community  (10:999) additive
    endif
  endif
end-policy
!
route-policy policy_5
  apply sample
  apply policy_3
end-policy
```

Given this sample configuration, the **show rpl unused as-path-set** command displays the following information:

```
RP/0/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
```

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 mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

Use the **show rpl unused community-set** command to display all the community sets that are not used in a policy at an attach point either directly or indirectly and are not referenced by any policies in the system.

Task ID

Task ID	Operations
route-policy	read

Examples

This example shows the following sample configuration:

```
router bgp 2
 address-family ipv4 unicast
 !
 neighbor 10.0.101.2
  remote-as 100
  address-family ipv4 unicast
   route-policy policy_1 in
  !
 !
 neighbor 10.0.101.3
  remote-as 12
  address-family ipv4 unicast
   route-policy policy_2 in
  !
 !
!

community-set comm_set_ex1
  65500:1,
  65500:2,
  65500:3
end-set
!
community-set comm_set_ex2
  65501:1,
  65501:2,
```

```
  65501:3
end-set
!
community-set comm_set_ex3
  65502:1,
  65502:2,
  65502:3
end-set
!
route-policy sample
  if (destination in sample) then
    drop
  endif
end-policy
!
route-policy policy_1
  if (destination in prefix_set_ex1) then
    set local-preference 100
  endif
  if (as-path in as_path_set_ex1) then
    set community  (10:333) additive
  endif
end-policy
!
route-policy policy_2
  if (destination in prefix_set_ex1) then
    if (community matches-any comm_set_ex1) then
      set community  (10:666) additive
    endif
    if (extcommunity matches-any ext_comm_set_rt_ex1) then
      set community  (10:999) additive
    endif
  endif
end-policy
!
route-policy policy_3
  if (destination in prefix_set_ex2) then
    set local-preference 100
  endif
  if (as-path in as_path_set_ex2) then
    set community  (10:333) additive
  endif
end-policy
!
route-policy policy_4
  if (destination in prefix_set_ex2) then
    if (community matches-any comm_set_ex2) then
      set community  (10:666) additive
    endif
    if (extcommunity matches-any ext_comm_set_rt_ex2) then
      set community  (10:999) additive
    endif
  endif
end-policy
!
route-policy policy_5
  apply sample
  apply policy_3
end-policy
```

Given this sample configuration, the **show rpl unused community-set** command displays the following information:

```
RP/0/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
```

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 mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

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
```

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 mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

Use the **show rpl unused prefix-set** command to display all prefix sets that are not used in a policy at an attach point either directly or indirectly and are not referenced by any policies in the system.

Task ID

Task ID	Operations
route-policy	read

Examples

This example shows the following sample configuration:

```
router bgp 2
 address-family ipv4 unicast
 !
 neighbor 10.0.101.2
  remote-as 100
  address-family ipv4 unicast
   route-policy policy_1 in
  !
 !
 neighbor 10.0.101.3
  remote-as 12
  address-family ipv4 unicast
   route-policy policy_2 in
  !
 !
!

prefix-set sample
  0.0.0.0/0,
  0.0.0.0/0 ge 25 le 32,
  10.0.0.0/8 ge 8 le 32,
  192.168.0.0/16 ge 16 le 32,
  224.0.0.0/20 ge 20 le 32,
  240.0.0.0/20 ge 20 le 32
end-set
!
```

```
prefix-set prefix_set_ex1
  10.0.0.0/16 ge 16 le 32,
  0.0.0.0/0 ge 25 le 32,
  0.0.0.0/0
end-set
!
prefix-set prefix_set_ex2
  220.220.220.0/24 ge 24 le 32,
  220.220.120.0/24 ge 24 le 32,
  220.220.130.0/24 ge 24 le 32
end-set
!
prefix-set prefix_set_ex3
  221.221.220.0/24 ge 24 le 32,
  221.221.120.0/24 ge 24 le 32,
  221.221.130.0/24 ge 24 le 32
end-set
!
route-policy sample
  if (destination in sample) then
    drop
  endif
end-policy
!
route-policy policy_1
  if (destination in prefix_set_ex1) then
    set local-preference 100
  endif
  if (as-path in as_path_set_ex1) then
    set community  (10:333) additive
  endif
end-policy
!
route-policy policy_2
  if (destination in prefix_set_ex1) then
    if (community matches-any comm_set_ex1) then
      set community  (10:666) additive
    endif
    if (extcommunity matches-any ext_comm_set_rt_ex1) then
      set community  (10:999) additive
    endif
  endif
end-policy
!
route-policy policy_3
  if (destination in prefix_set_ex2) then
    set local-preference 100
  endif
  if (as-path in as_path_set_ex2) then
    set community  (10:333) additive
  endif
end-policy
!
route-policy policy_4
  if (destination in prefix_set_ex2) then
    if (community matches-any comm_set_ex2) then
      set community  (10:666) additive
    endif
    if (extcommunity matches-any ext_comm_set_rt_ex2) then
      set community  (10:999) additive
    endif
  endif
end-policy
!
```

```
route-policy policy_5
  apply sample
  apply policy_3
end-policy
--------------------------
ext_comm_set_ex3
```

Given this sample configuration, the **show rpl unused prefix-set** command displays the following information:

```
RP/0/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
```

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 mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

Use the **show rpl unused rd-set** command to display all of the RD sets that are not used in a policy at an attach point either directly or indirectly and are not referenced by any policies in the system.

Task ID

Task ID	Operations
route-policy	read

Examples

The **show rpl unused rd-set** command displays the following information:

```
RP/0/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.
```

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 mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

Use the **show rpl unused route-policy** command to display route policies that are defined but not used at an attach point or referenced from another policy using an **apply** statement.

Task ID

Task ID	Operations
route-policy	read

Examples

This example shows the following sample configuration:

```
RP/0/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
```

transmit (ePBR)

The command forwards packets through the normal IP routing table lookup if they match the associated class map criteria configured under a PBR class within a policy map.

transmit

Command Default

None.

Command Modes

Policy map configuration mode.

Command History

Release	Modification
Release 25.3.1	This command was introduced.

Task ID

Task ID	Operation
qos	read, write

```
Router(config)#policy-map type pbr INGRESS-SECURITY-POLICY
Router(config-pmap)#class type traffic MANAGEMENT-TRAFFIC
Router(config-pmap-c)#transmit
Router(config-pmap-c)#exit
```



Static Routing

- address-family (static), on page 782
- address-family ipv4 unicast (static), on page 783
- address-family ipv6 unicast (static), on page 784
- maximum path (static), on page 785
- router static address-family ipv4 unicast, on page 786
- route (static), on page 787
- router static, on page 790
- show static interfaces, on page 791
- show static vrf, on page 792
- vrf (static), on page 793

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

Usage Guidelines

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.

Examples

The following example shows how to enter IPv6 unicast address family mode:

```
Router(config)# router static
Router(config-static)# address-family ipv6 unicast
Router(config-static-afi)#
```

Related Commands

Command	Description
Route Static	Establishes a static route.

address-family ipv4 unicast (static)

To enter IPv4 address family configuration mode while configuring static routes, use the **address-family ipv4 unicast** command in the appropriate configuration mode. To disable support for IPv4 address family, use the **no** form of this command.

address-family ipv4 unicast *prefix mask* **remote-next-hop** *ip4-address* **tunnel** *gue-v1* **source-ip** *ip-address* [**dscp** *value*] [**ttl** *value*]

Syntax Description

prefix mask	Specifies the destination prefix (A.B.C.D/length or X:X::X/length).
remote-next-hop	Specifies the remote tunnel IP4 endpoint address in default VRF.
tunnel *gue-v1*	Specifies the encapsulation type.
source-ip *ip-address*	Specifies the IPv4 source IP address.
dscp *value*	(Optional) Specified the DSCP value to be set in GUE encapsulation IP header. Ranges between 0 and 63.
ttl *value*	(Optional) TTL value to be set in GUE encapsulation IP header. Ranges between 0 and 255.

Table 132: Command History:

Release	Modification
Release 25.4.1	This command was modified to include these keywords: **remote-next-hop**, **tunnel**, **source-ip**, **dscp**, and **ttl**.
Release 7.0.14	This command was introduced.

Command Modes

Router static configuration

Usage Guidelines

IPv4 option in **source-ip** is visible only for IPv4 next-hop.

Examples

This example shows how to configure IPv4 GUEv1 static tunnel:

```
  Router(config)# router static
  Router(config-static)# address-family ipv4 unicast
  Router(config-static-afi)# 9.8.7.6/32 remote-next-hop 5.5.5.6 tunnel gue-v1 source-ip
1.1.1.1 dscp 46 ttl 250
```

Related Commands

Command	Description
Route Static	Establishes a static route.

address-family ipv6 unicast (static)

To enter IPv6 address family configuration mode while configuring static routes, use the **address-family ipv6 unicast** command in the appropriate configuration mode. To disable support for IPv6 address family, use the **no** form of this command.

address-family ipv6 unicast *prefix mask* **remote-next-hop** *ip4-address* **tunnel** *gue-v1* **source-ip** *ip-address* [**dscp** *value*] [**ttl** *value*]

Syntax Description

prefix mask	Specifies the destination prefix (A.B.C.D/length or X:X::X/length).
remote-next-hop	Specifies the remote tunnel IPv6 endpoint address in default VRF.
tunnel *gue-v1*	Specifies the encapsulation type.
source-ip *ip-address*	Specifies the IPv6 source IP address.
dscp *value*	(Optional) Specifies the DSCP value to be set in GUE encapsulation IP header. Ranges between 0 and 63.
ttl *value*	(Optional) Specifies the TTL value to be set in GUE encapsulation IP header. Ranges between 0 and 255.

Table 133: Command History:

Release	Modification
Release 25.4.1	This command was modified to include these keywords: **remote-next-hop**, **tunnel**, **source-ip**, **dscp**, and **ttl**.
Release 7.0.14	This command was introduced.

Command Modes

Router static configuration

Usage Guidelines

IPv6 option in **source-ip** is visible only for IPv6 next-hop.

Examples

This example shows how to configure IPv6 GUEv1 static tunnel:

```
  Router(config)# router static
  Router(config-static)# address-family ipv6 unicast
  CPU0:router(config-static-afi)# 11:1:1::1/128 remote-next-hop 10.10.10.11 tunnel gue-v1
source-ip 2.2.2.2 dscp 8 ttl 250
```

Related Commands

Command	Description
Route Static	Establishes a static route.

maximum path (static)

To change the maximum number of allowable static routes, use the **maximum path** command in static router configuration mode. To remove the **maximum path** command from the configuration file and restore the system to its default condition, use the **no** form of this command.

maximum path { **ipv4** | **ipv6** } *value*
no maximum path { **ipv4** | **ipv6** } *value*

Syntax Description

ipv4 | **ipv6**	Specifies IP Version 4 (IPv4) or IP Version 6 (IPv6) address prefixes.
value	Maximum number of static routes for the given AFI. The range is 1 to 140000.

Command Default

value : 4000

Command Modes

Static router configuration

Usage Guidelines

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.

Examples

The following example shows how to set the maximum number of static IPv4 routes to 100000:

```
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:

```
Router(config-static)# no maximum path ipv4 100000
```

Related Commands

Command	Description
route static	Enters static router configuration mode.
show route static	Displays the static routes in a routing table.

router static address-family ipv4 unicast

Use this command to allow an IPv6 address to be entered as a next-hop for an IPv4 prefix.

router static address-family ipv4 unicast *ipv4-prefix/mask interface ipv6-next-hop*

Syntax Description

ipv4-prefix mask	Specifies the IPv4 destination prefix (A.B.C.D/length or X:X::X/length).
interface	Specifies the interface that must be paired with an IPv6 next-hop address.
ipv6-next-hop	Specifies the IPv6 gateway address of the next router in the path toward the destination.

Table 134: Command History:

Release	Modification
Release 26.1.1	This command was introduced.

Command Default

All static routes belong to the default VRF in the address family configuration mode.

Command Modes

Router static configuration

Examples

This example shows how to configure recursive path. Use this option when the IPv6 next-hop is reachable via another route in the routing table.

```
  Router(config)# router static
  Router(config-static)# address-family ipv4 unicast
  Router(config-static-afi)# 1.1.0.0/16 1:1:1::1
Router(config-static-afi)# commit
```

This example shows how to configure IPv6 next-hop via an interface . Use this option to specify exactly which interface must be used to reach the IPv6 next-hop.

```
  Router(config)# router static
  Router(config-static)# address-family ipv4 unicast
  Router(config-static-afi)# 1.1.0.0/16 GigabitEthernet0/0/0/0 1:1:1::1
Router(config-static-afi)# commit
```

Related Commands

Command	Description
Route Static	Establishes a static route.

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*] [**behavior** | *utef*] [**track** *track-object-name*] [**tunnel-id**] [**vrflabel** *vrf-label*] [*distance*] [**description** *text*] [**tag** *tag*] [**permanent**] }

Syntax Description

prefix / *mask*	IP route prefix and prefix mask for the destination. The network mask can be specified in either of two ways: • The network mask can be a four-part, dotted-decimal address. For example, 255.0.0.0 indicates that each bit equal to 1 means the corresponding address bit is a network address. • The network mask can be indicated as a slash (/) and number. For example, /8 indicates that the first 8 bits of the mask are 1s, and the corresponding bits of the address are the network address.
ip-address	IP address of the next hop that can be used to reach that network. • For IPv4 address–the IP address is required, not optional, if the interface type and interface-path-id arguments are not specified. You can specify an IP address and an interface type and interface path. • For IPv6 link-local address–the interface type and interface-path-id arguments are required. The route is not valid, if the interface type and interface-path-id arguments are not specified. **Note** A forwarding router's IP address or an interface or virtual interface path ID can be configured, in any order.
behavior	Define the behavior by selecting the appropriate option.
utef	Static programs the utef SID with the context of controller address.
type	Interface type. For more information, use the question mark (?) online help function.
interface-path-id	Physical interface or virtual interface. **Note** Use the **show interfaces** command to see a list of all interfaces currently configured on the router. For more information about the syntax for the router, use the question mark (**?**) online help function. **Note** A forwarding router's IP address or an interface or virtual interface path ID can be configured, in any order.

distance	(Optional) Administrative distance. Range is 1 to 254.
description *text*	(Optional) Specifies a description of the static route.
tag *tag*	(Optional) Specifies a tag value that can be used as a match for controlling redistribution using route policies. Range is 1 to 4294967295.
permanent	(Optional) Specifies that the route is not removed from the routing table, even if the next-hop interface shuts down or next-hop IP address is not reachable.
track *track-object-name*	Installs static routes in RIB only if track is UP.
tunnel-id *tunnel-id*	Specifies a Tunnel ID.
vrflabel *vrf-label*	Specifies a VRF label.

Command Default

No static route is established.

Command Modes

Address family configuration

Table 135: Command History:

Release	Modification
Release 24.1.1	The *utef* keyword was introduced.
Release 7.0.14	This command was introduced.
Release 7.3.15	The conditional installation of static route in routing table (RIB) based on object state functionality has been enhanced to use event driven notification, instead of the polling mechanism used earlier.

Usage Guidelines

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.

A static route is always associated with a VPN routing and forwarding (VRF) instance. The VRF can be the default VRF or a specified VRF. Specifying a VRF allows you to enter VRF configuration mode where you can configure a static route. If you do not specify a VRF you can configure a default VRF static route.

Use the **router static** command to configure static routes. To configure a static route, you must enter router static configuration mode and then enter an address family configuration mode or VRF configuration mode. See the **vrf (static)** command for information on configuring a static route in VRF configuration mode. After

you enter an address family mode, you can enter multiple static routes. The following example shows how to configure multiple static routes in IPv4 and IPv6 address family configuration modes:

Note You cannot create a VRF named default, but you can reference the default VRF.

Note The conditional installation of static route in routing table (RIB) based on object state functionality has been enhanced to use event driven notification, instead of the polling mechanism used earlier. This event driven approach allows faster detection and notification of object state changes.

The following example shows how to configure the Sink node:

```
Router(config)#router static
Router(config-static)# address-family ipv6 unicast
Router(config-static-afi)#fccc:cc00:9000:fef1::/64 segment-routing srv6 endpoint behavior
utef controller-address fccc:cc00:7::
```

Examples

The following example shows how to configure IPv6 unicast address family static routes:

```
Router(config)# router static
Router(config-static)# address-family ipv6 unicast
Router(config-static-afi)# 2b11::327a:7b00/120 GigabitEthernet0/2/0/7
Router(config-static-afi)# 2b11::327a:7b00/120 GigabitEthernet0/6/0/0
Router(config-static-afi)# 2b11::327a:7b00/120 2b11::2f01:4c
Router(config-static-afi)# 2b11::327a:7b00/120 2b11::2f01:4d
Router(config-static-afi)# 2b11::327a:7b00/120 2b11::2f01:4e
Router(config-static-afi)# 2b11::327a:7b00/120 2b11::2f01:4f
Router(config-static-afi)# 2b11::327a:7b00/120 2b11::2f01:50
```

router static

To enter static router configuration mode, use the **router static** command in mode. To remove all static route configurations and terminate the static routing process, use the **no** form of this command.

router static
no router static

Syntax Description

This command has no arguments or keywords.

Command Default

No static routing process is enabled.

Command Modes

Examples

The following example shows how to enter static router configuration mode:

```
Router(config)# router static
Router(config-static)#
```

Related Commands

Command	Description
Address family static	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	Establishes a static route.

show static interfaces

Use this command to display a specific notification indicating if an interface is being tracked for forwarding cross-afi IPv6 traffic.

show static interfaces

Table 136: Command History:

Release	Modification
Release 26.1.1	This command was introduced.

Command Default

All static routes belong to the default VRF in the address family configuration mode.

Command Modes

Router static configuration

Usage Guidelines

No specific guidelines impact the use of this command.

Examples

This example displays the interface state is **Up** and tracked for cross-afi traffic.

```
Router# show static interfaces
```

show static vrf

Use this command to display IPv6 next-hops and the topology table that lists IPv4 prefixes.

show static vrf default ipv4 *next-hops topology*

Syntax Description

next-hop	Displays IPv6 next-hops and their resolution state, Resolved or Unresolved specifically for the IPv4 routing process.
topology	Displays the topology table, listing IPv4 prefixes alongside their corresponding IPv6 next-hops and interfaces.

Table 137: Command History:

Release	Modification
Release 26.1.1	This command was introduced.

Command Default

Not applicable.

Command Modes

Router static configuration

Usage Guidelines

No specific guidelines impact the use of this command.

Examples

This example shows how to verify if the static routing process has successfully resolved the IPv6 next-hop. Look for the status **Resolved** next to the IPv6 address.

```
Router# show static vrf default ipv4 next-hops
```

This example shows how tcheck the Routing Table (RIB) and verify that the IPv4 prefix is present and shows the IPv6 next-hop.

```
Router# show route ipv4 unicast 1.1.0.0/16
```

vrf (static)

To configure a VPN routing and forwarding (VRF) instance and enter VRF configuration mode, use the **vrf** command in router configuration mode. To remove the VRF instance from the configuration file and restore the system to its default condition, use the **no** form of this command.

vrf *vrf-name* **vrf RED urpf-lookup-ipv4** | **urpf-lookup-ipv6**
no vrf *vrf-name*

Syntax Description

vrf-name	Name of the VRF instance. The following names cannot be used: all, default, and global.
vrf RED urpf-lookup-ipv4	Introduced in IOS XR Release 26.2.1, it enables the IPv4 URPF lookup bit for the VRF.
vrf RED urpf-lookup-ipv6	Introduced in IOS XR Release 26.2.1, it enables the IPv6 URPF lookup bit for the VRF.

Command Default

No default behavior or values

Command Modes

Static router configuration

Usage Guidelines

Use the **vrf** command to configure a VRF instance. A VRF instance is a collection of VPN routing and forwarding tables maintained at the provider edge (PE) router.

A static route is always associated with a VRF, which is entirely user configurable. Static route is unique within a VRF. A static route can point to a next-hop interface, next-hop IP address, or both, which can be resided in the same VRF configured for the static route or in a different VRF. For example, routes 172.168.40.0/24 and 172.168.50.0/24 are configured as follows:

```
router static
vrf vrf_A
address ipv4 unicast
172.168.40.0/24 loopback 1
172.168.50.0/24 vrf vrf_B 192.168.1.2
```

Routes 172.168.40.0/24 and 172.168.50.0/24 belong to vrf_A. Route 172.168.50.0/24 is not installed in vrf_A until next-hop 192.168.1.2 (a vrf_B route) is reachable.

If you are configuring a default VRF route, you do not need to enter VRF configuration mode. For example, routes 192.168.1.0/24 and 192.168.2.0/24 are configured as follows:

```
router static
 address ipv4 unicast
  192.168.1.0/24 loopback 5
  192.168.2.0/24 10.1.1.1
```

Routes 192.168.1.0/24 and 192.168.2.0/24 are default VRF routes.

Note You cannot create a VRF named default, but you can reference the default VRF.

Examples

The following example shows how to configure a VRF instance and enter VRF configuration mode:

```
Router(config)# router static
Router(config-static)# vrf vrf-1
Router(config-static-vrf)#
```

This example shows how to configure URPF lookup on a specific VRF:

```
Router# configure
Router(config)# vrf RED
Router(config-vrf)# urpf-lookup-ipv4
Router(config-vrf)# urpf-lookup-ipv6
Router(config-vrf)# commit
```

Related Commands

Command	Description
Address family static	Enters address family configuration mode and allows you to configure a static route.